

VOLUME 14 PART 1 APRIL 1995

BRITISH TELECOMMUNICATIONS ENGINEERING

Included in this Issue

BT Networks and Systems

*Eliminating the Handicap of
Special Needs*

*Network Measurement and
Performance*

When I'm 64



**The Journal of The Institution of
British Telecommunications Engineers**



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 14 ■ PART 1 ■ APRIL 1995

Guest Editorial: Shaping BT for the Future Alan Rudge	1
BT Networks and Systems—A Platform for Success Chris Earnshaw	2
<i>Making it Happen</i> Network Measurement and Performance Deryck Rogers and David Hand	5
<i>Human Factors</i> Eliminating the Handicap of Special Needs Stephen Furner and Martin Cooper	12
<i>Telecommunications in the 21st Century</i> Challenges in Information Visualisation Graham Walker	17
When I'm 64—Technology and Telecommunications Beyond 2025 Ian Pearson	26
Aspects of Optical Transparency Peter Cochrane and David Heatley	33
<i>Mobile and Radio Communications</i> Global System for Mobile Communications—The Development of the GSM Standard Peter M. Adams and Nigel Wall	38
The St. Albans Study John Prior and Kathy Chaplin	46
Developing Sales Engineering into Engineering Consultancy Andy Jeffery and Colin Banks	50
Telecommunications Fraud Management—Reducing Costs and Winning Business James Messham, Phil Kelly and Dean Smith	58
High-Voltage Protection of Telecommunication Circuits John Gerrard and Walter Sweetenham	63
The Development of Power and Building Engineering Services in Telecommunications John Parsons	71
TRIBUNE—A Broadband ISDN Test Bed Chris Shepherd	81
<i>Regular Features</i> BT News	89
Industry News	91
Book Reviews	95
Notes and Comments	96
IBTE Contacts	Inside back cover

Board of Editors

Chairman John Prior, M.Sc.

Managing Editor Paul Nichols, B.Sc. **Secretary** Stephen Wilson, B.Sc. **Treasurer** Chris Fennell, B.A., F.C.M.A.

Assistant Editor Richard Kemp, B.Sc., C.Eng., M.I.E.E.; Alan Bealby, M.Sc., F.I.E.E.; Richard Beet, M.A.I.E.(DIP.), F.R.A.S., F.B.I.S.; Martin Evans, B.Sc., Ph.D.; Warwick Flury; Fred Landery, B.A.; Sam Quinn, B.A., DIP.M.; Kevin Rose, B.Sc., M.Sc., C.Eng., M.I.E.E., M.B.C.S.; Murray Trantor, B.Sc., C.Eng., F.I.E.E., D.M.S., M.B.I.M.; Bob Ward, B.Sc., C.Eng., M.I.E.E., M.I.MGT.; Keith Ward, C.Eng., F.I.E.E., M.B.I.M.

Theme Editors

Making it Happen Alan Bealby

Mobile and Radio Communications Ian Groves and Nigel Wall

Telecommunications in the 21st Century Peter Cochrane and Bill Whyte

Human Factors Martin Cooper

Published by: The Institution of British Telecommunications Engineers, 2-12 Gresham Street, London EC2V 7AG

ISSN 0262-401X

Alan Rudge

Shaping BT for the Future



Underpinning the reorganisation and realignment of the BT Group which took effect on 3 April 1995 are a number of strong underlying principles. The first is an acknowledgement of the *technological convergence* which is occurring between computing and telecommunications. This convergence was predicted long ago but has taken a significant time to penetrate all aspects of our business.

Last year we made the decision to merge our information technology and telecommunications research and development activities. This decision was driven by the recognition of the similarities in the expertise and procedures necessary in the development of complex software systems, whether they be directly for telecommunications applications or for the information systems which support their operation.

Many benefits became apparent following this merger, not least of which were advantages which arose from better teamworking in the development of highly integrated systems, which are becoming the norm in BT. It became increasingly evident that the convergence was not only in the technology. There was also a convergence in the skills of our

people in computing systems and networks which was creating a new 'core competence' for BT.

The benefits achieved in Development provided strong support for BT's 'Breakout' proposal for extending this integrated approach to all of our operations in networks and computer systems. This experience helped pave the way for the decision to bring together our networks and computer systems capabilities into one Division, with a single *design, plan and build function* for both activities.

There is little doubt that if we are to meet the needs of the *Information Age* in the 21st century, then we can no longer think of our networks as combinations of switches and cables, supported by some computer systems. A better conceptual model of our network is that of a *giant distributed computer*, which demands the complete integration of the transport functions (the cables and switches) with the intelligence, management and control, and the other operating systems provided by our computing functions. This integrated approach is absolutely essential if we are to provide the kind of responsive, flexible and efficient services which our customers will undoubtedly expect.

BT's core business activities have been structured functionally for some years now. The recent reorganisation strengthens this approach and endorses the need to enhance our capabilities in managing cross-divisional activities by using matrix-management techniques. The organisation structure is arranged along the lines of *requirements-delivery-coordination* principles which reinforces the customer-supplier culture within the business as well as externally.

While matrix-working across the Divisions will be the norm, in general the customer-facing Divisions will lead

on defining customer *requirements*, the Networks and Systems Division will lead on providing the engineering solutions and their *delivery*, while the newly created Group Business Management Division will provide any necessary *coordination* and prioritisation.

My own role in the new structure is to share the load of the Chairman and the Group Managing Director in moving BT forward on the many fronts in which we are currently engaged. For example, my new activities include taking particular responsibility for major group programmes, such as our activities in multimedia and the evolution of our network into a full broadband capability. We shall also be extending matrix-management practices across the business and I shall be overseeing the implementation of the new management accounting and information systems which are essential to efficient matrix-working.

I am also responsible for the completion of BT's 'Breakout' process re-engineering programme and the realisation of the benefits arising from it. I shall remain the Group's senior technology spokesman and will retain responsibility for the corporate research budget and chairmanship of the Technology and Development Committee. With my other activities on the Main Board, the Chairman's Advisory Group, the Investment Committee and the Group MD's Committee and Quality Council, I do not expect to be bored for the foreseeable future! There is much to be done and I look forward to helping to shape BT for what will clearly be an exciting and dynamic future.

Alan Rudge

Deputy Group Managing Director, BT

BT Networks and Systems— A Platform for Success



*Chris Earnshaw, Managing Director,
BT Networks and Systems*

*Our objectives are
to give highest
priority to
optimising
performance of
the BT Group as a
whole and to
enable the delivery
of the best value
and performance
to BT's customers.*

Background

As we move towards the 21st century, success for BT will depend critically on our ability to respond to changing demands of the UK and global market and to establish clear competitive advantage, as the industry makes the transition from 'telephony' to 'multimedia and information services'. The recently announced Networks and Systems organisation has a vital role to play in creating BT's 'Platform for Success'.

The formation of the new organisation is specifically aimed at helping us to manage the transition to our future business by enabling the closer integration of the previously separate worlds of networks and computer systems—at a design, development and operational level. Over recent years, the network, and the computer systems which support its real-time operation, have grown in complexity to form a vast interconnected web. The coherent real-time performance of the total system is now critical to the delivery of service to meet the increasing needs and expectations of customers. Although our organisation had undergone many changes in recent years, there were many areas where the structure still reflected the operational and technology boundaries of previous decades, prior to the creation of a fully digital and highly automated infrastructure. Our new organisational structure will align our activities around coherent end-to-end processes and take us an important step forward from the creation of Customer Field Services (UKCFS) and the integration of the core and global networks.


Together with the customer-facing divisions and Group Business

Management, Networks and Systems forms the core of BT's overall business. Our objectives are to give highest priority to optimising performance of the BT Group as a whole and to enable the delivery of the best value and performance to BT's customers. We will establish relationships with the customer-facing divisions, which reflect their priorities for developing BT's business in the marketplace.

Facing the Future

Over the past 18 months or so, a fundamental reappraisal of our operational processes has been undertaken through Project Breakout. This has involved both critical self-analysis against best practice and some fundamental rethinking about how to simplify and streamline our operations to exploit more effectively the many billions of pounds we have invested in recent years. In many cases, in line with the experience of other companies, we find that we have added complexity, cost and delay into our operations without seizing the opportunity for more radical transformation that new technologies potentially allow. We have also been able to judge critically our performance and targets against those of other leading operators and emerging new players who will be our competitors in the new information age.

By integrating the key operational units with the professional systems planning and development resources, we will considerably simplify our activities. We will be better placed to realise the potential of the emerging intelligent and broadband networks, streamline the evolution of our



products and services, focus our operational support and information systems on real-time customer services, and achieve world-class performance in effectiveness and efficiency.

We will work as a team, partnering with each other and the rest of BT, to deliver maximum value to our customers. We will look to our managers to exhibit a 'bias for action' and be role models of professionalism and coaches for their teams and effective in identifying and removing barriers to their progress. Other team members will assume more responsibility, as we push decision-making closer to the customer, and embrace change as we continue to improve the way we work. A passion for excellence and drive to achieve world-class standards must pervade all we do.

Realigning Our Organisation

Network and Service Operations

The role of Network and Service Operations division is to deliver world-class service by running the world's most effective telecommunications platforms. Key measures of performance will be quality of services to customers (defined as consistently meeting agreed performance levels), and efficiency (defined as world-class financial productivity). We can think of the division as the

'production line of BT'. As well as seeking economies of scale, we must support increasing levels of differentiation and an expanding portfolio of services.

Our operations will continue to be dominated by provision and repair of local access for telephony services. It is in this area we must take major steps to achieve significant improvements in quality and service standards.

We are also seeing the emergence of voice, data and other broadband services on a more significant scale, and our operational processes and control systems need to evolve to cater for these to ensure we can deliver integrated services that our customers expect, as communication assumes a greater role in both personal lifestyles and business operations.

Design and Build

BT's marketplace is rapidly evolving both in scope, geography and competition. Our operational platform must respond to these challenges. Design and Build will be responsible for creating a coherent and rapidly evolving systems platform to support the delivery of service and the operational capabilities to support customers. Effective investment in appropriate technologies will be critical in ensuring our platform continues to provide competitive capabilities and performance. Speed to market, as well as cost performance, are critical success factors.

We will examine the links between professional engineering design and field construction teams to ensure we have efficient implementation plans, which balance the need for proactive investment in new capacity and technologies with the ongoing reactive work to maintain standards of service. The transformation of the UK access network will continue to be a high priority.

The creation of an intelligent platform to support an increasing range of broadband and narrowband applications will be high on our agenda. In doing this, we must strive for maximum synergies between our UK platform and those we are creating in other countries, as well as through Concert, our joint venture with MCI, focused on global services.

Development

The timely availability of systems and new technologies will be vital in maintaining our competitive advantage. The Development division will provide a key source of network, service and application technologies. While we continue to develop our own areas of excellence, it will become increasingly important that we have the ability to integrate these with externally sourced technologies and systems. Competitive advantage lies not in the component systems themselves, but our ability to produce differentiated solutions in terms of cost, quality and capability. Our investment in new systems, services

and application innovations provide an important means for continuing to grow our business.

Supply Management

Almost all aspects of the business are dependent upon the contribution of external suppliers. The Supply Management division will ensure that we establish and operate effective supplier relationships that deliver materials and services, systems and software that contribute directly to improving BT's competitive advantage and increasing its profitability. Optimising each supply chain from suppliers to ultimate customer offers significant scope for continuous improvement based on 'win-win' opportunities for everyone involved.

Supporting Our Operations

Process Control and Measurement

We have established a Process Control and Measurement division to develop coherent end-to-end operational process improvements which attain world-class standards of performance. The emphasis on measurement is deliberate to ensure we drive our programmes for continuous business improvement through appropriate operational measures. The division will play a key role in the allocation of resources to drive for performance against predetermined standards.

Finance

Effective management of our financial resources will be critical to meeting customers' expectations and providing value for money in an increasingly competitive environment.

We will continue to refine our fiscal and planning process to align with best commercial practice and create an environment that drives for continuous improvement and exploitation of the platform infrastructure. This will include ensuring that the costs of services and resources are

recognised and borne by the appropriate customer facing division.

Leadership and people

As a business unit at the heart of the convergence between technologies and new markets, we will see continuing change. The ability and skills of our people to manage and respond to these changes will be critical to our success. The support and development of our people is a major element of our management challenge.

Our objective is to ensure our workforce is recognised as world class in its field of activity. The emphasis will shift from simply reducing numbers to reducing our total costs to deliver competitively priced services. This will involve continuous rebalancing of skills and simplification of work practices through effective use of technologies and information systems.

Success

Success will be achieved when three conditions are met. First, BT's customers must be delighted with the value and performance of the services they receive from Networks and Systems as part of the greater BT team. Second, our internal customers must be likewise pleased with our ability to respond to their changing needs. Finally, our people must feel fulfilled by the role they play and the work they do.

Networks and Systems has a vital role to play in ensuring success for BT's customers, shareholders and people as we move towards the next millennium.

Biography

Chris Earnshaw is Managing Director of Networks and Systems for BT and member of the Group Managing Director's Committee and Quality Council, GQC. He is responsible for the planning and operations of the company's infrastructure and service delivery. This includes the company's network, supporting information systems, research and development, as well as the corporate procurement programme.

Prior to this new role, he was the President and Chief Executive Officer of Concert, the BT/MCI joint venture company launched in 1994, having been closely involved in the formulation of the concepts behind the joint venture over the previous year. Headquartered in Virginia, USA, Concert was established to provide business services for companies operating on a global basis.

Before moving to the USA, Chris had been Managing Director of BT's Worldwide Networks division, responsible for BT's UK and overseas networks, since 1991.

Chris originally joined BT in 1972 and was involved in a variety of assignments related to the design and implementation of digital network services, including the System X development programme, modernisation of the trunk network and introduction of 0800. In 1989, he was appointed Principal Director of BT (UK) and member of the BT (UK) Management Board as Director Network, with responsibility for the design and investment programmes for the UK network.

He is a Director of the International Engineering Consortium based in the US, and until 1994 was a Director of the Cellnet Group. He is a graduate of Sheffield University and Fellow of the Institution of Electrical Engineers.

Deryck Rogers and David Hand

Network Measurement and Performance

In order to be able to improve the performance of a telecommunications network, it is important to know how well it is working now. Previously this was achieved through making test calls on the network, but this could only check the performance of a small part of the network. This article explains how 'live calls' can be sampled to provide a more representative analysis of the performance of the network.

Introduction

For many years one of the basic measures of performance, call failures, has been derived by analysing the results from large volumes of test calls. Within BT since 1979, measurement and analysis centres (MACs)¹ held the systems delivering those measures. During 1988 a study was undertaken to identify a replacement for MAC, which would reach the end of its life expectancy in the mid 1990s. Not only was there an ongoing need to measure the performance of the basic telephony call, there was also a growing need to address the performance of new products and services, which were being delivered to what would become a highly competitive market. Cost reduction was another key consideration so any new system must also be more cost effective than the system it was to replace. Hence, the requirement was to measure more for less. The introduction of the new system has challenged the widely accepted view, throughout the world, that only test calls can be used to derive the measure of call performance. Rather than attempting to mimic the customer in calling patterns, the proposal was to use actual calls, to monitor their progress, and derive measures based on the actual experienced performance. The rollout of the replacement system has just been completed, and this article describes the new system, highlights some of the different views of the network this has provided and gives examples of the unexpected benefits this approach has provided. Before describing the new system, there is a need to understand some network performance terms and the historical

context of network performance measurement.

What is Network Performance?

A number of previous papers have established the basis of call performance measures using the various call phases^{2,3}.

Within any switched network the call phases can be summarised into four areas:

- initial access to the network,
- connection establishment and retention,
- information exchange, and
- clear down, which includes the raising of a charge, where appropriate.

Each of these phases has an associated set of measures (Table 1).

Clearly, all these phases come into play sequentially for an ordinary point-to-point telephone call, but the introduction of new network services complicates this progression through the phases. Three-way calling, for example, disrupts the previously straightforward situation, where a second call attempt is initiated before the first connection has been cleared. This is not the only complication for a test-call sending system, which traditionally has been a measure of the technology, not of what the customer experiences. This is discussed in detail in Anthony Oodan's work in the IBTE *Structured Information Programme*². Suffice to say that today both the measurement of the performance of the technology and

Table 1

Call Phase	Description	Typical Measures
Initial access	The ability to access the desired network or service.	Availability Downtime Mean time between failures
Call set-up	The ability to send a valid address, for the network to connect to that valid address, and to sustain that connection for the call duration.	Call failures
Information exchange	The ability to transfer information (speech/data) across the already established connection without obstruction or error.	Call clarity: —noise —echo, etc. Error rates: —errored seconds
Clear down	The ability to release all equipment used in the connection and return them to the common pool for re-assignment. Includes the release of both called and calling equipments.	As for set-up (A separate measure is rarely needed.)

Table 2

Call Sequence	Number of Test Calls
Own exchange calls	1000
Local call area	800
'National' calls	500
Total monthly calls	2300

Introduced in the late 1970s, MAC comprised some 63 GEC20/50 processors with ancillary calling and answering equipment, located throughout the country. Not all exchanges were equipped with MAC calling equipment, the criterion being those exchanges serving 1000 or more customers. MAC covered some 500 exchanges serving approximately 80% of all customers. Test call results were analysed into the, now well-known, plant engaged (PE) or plant defect (PD) categories. In an attempt to make the results more representative, analysis of customers' actual calling patterns was regularly undertaken and used to 'steer' the test calls to similar destinations in similar proportions. Results were also weighted to reflect the significance of different sizes of exchange.

Fixed numbers of test calls were originated from each exchange every month in specific calling patterns as shown in Table 2.

Approximately 2 million test calls per month were made across the network and, from these, performance reported, blackspots identified and subsequent remedial action plotted. For many years MAC was the source of the 'percentage of call failures due to BT' reported in the BT Annual Report and Accounts, and to OFTEL. The use of MAC has been a significant success, both as a network performance measurement tool and as a device to aid maintenance staff to locate problems. Figure 1 shows the change in performance since 1990, which has been achieved through much better direction of maintenance effort, and latterly through maintaining that emphasis during modernisation. Despite the success of MACs in the past, there were a number of reasons that made it (and any other test-call sending systems) unsuitable in a modern telecommunications network. Following an extensive




Figure 1—UK average call failure rates from October 1990

the quantification of customer experiences are both valid views.

Historic Perspective—Test Call Senders

The traditional approach to monitoring call performance in an essentially plain old telephone system (POTS) environment was straightforward. Customer behaviour was characterised as either making local or national calls, and by using test calls it would be possible to imitate their behaviour to derive a measure of the network's performance. Test calls could be literally someone making a

test call and logging the result, but by the late 1970s a number of automatic test call senders became available. For an automatic assessment to be possible, it is necessary to have not only some equipment capable of initiating calls, but also a remote answering capability to confirm that a correct connection had been established. This places a restriction on the destinations that could be measured. For as long as customer behaviour conformed to the POTS view, this was not a problem and one particular test-call sending system, the MAC, has been a conspicuous success.

study, it was decided to use live call sampling as the basis of call performance measurement in the modern switched network.

Live Call Sampling

Most modern exchange systems have a live-call sampling facility as part of their management statistics package, supplied either as a standard feature or as an optional extra. Both System X and AXE10, as supplied to BT, have the facility. The detailed operation of the facility varies with each system, but the general approach is sufficiently generic to fit into a single overview description. (While the enhanced TXE4 also has the facility, the data derived from it is unsuitable for end-to-end performance measurement purposes.)

Overview of sampling

All calls originating, terminating or passing through digital exchanges are eligible to be sampled. Once the sampling system has been activated and the sampling rate (n) chosen, the exchange software counts the call attempts until n is reached, at which point, the n th call is flagged for analysis. Throughout the setup, information transfer and clear-down of that call, specific information will be recorded and at the end of that call the information will be written to disk for subsequent onward transmission to the processing system, where the data is analysed and results produced. At the time when the n th call is first flagged, no information is known about the intended call, except that a call attempt is in progress. The attempt may not mature, and may be abandoned before dialling (for example, accidentally disturbing a handset, and its subsequent replacement, will be seen as a call attempt by the switch). Similarly at that time there is no way of telling if the intended call will be a local, national or international call, if the call will remain in the host's network, or if it will go to a competitor's network. All that is known is the date and time of

the attempt, and some information about which equipment has been used in the call request. There is no information recorded that will allow the identification of the calling customer, although general information such as the tariff group (Business or Residential) will be recorded.

As the call attempt progresses, the information is updated as more information becomes available. Dialed digits, and identification of the routing the call has used within the exchange will all be added. One other crucial piece of data is also logged, and is constantly updated during the call setup, namely what can be generically termed the 'failure reason', which might better be described as the reason for call termination. Should any call terminate, either in setup, during information transfer, or clear-down, then the termination reason will provide a specific indication of what happened. When used in association with the call phase it is possible to attribute blame for the failure. There are many hundreds of different termination reasons, identifying many different eventualities. These need careful interpretation before meaningful results can be produced. It must be recognised that every call terminates at some time, so every call has a failure reason appended. This information is derived from the switch as it processes the call setup, but will also derive information from remote switches over the C7 signalling channels. This makes it possible to monitor the progress of the call on almost an end-to-end basis. Of course, most calls have a termination reason that shows that they have progressed through all the call phases, to a normal termination—in other words, successful calls.

Despite the large number of failure categories, there are some failures that cannot be identified. While there can be no monitoring of the originator's calling loop, or any event before the call initiation signal is recognised by the exchange, some failures in the called customer's local

loop and own equipment can be recognised. Call clarity cannot be measured by this system, nor can mis-routing, although it could be argued that poor performance in these areas will result in very short call holding time. The great strength of this measurement system is that it measures calls going anywhere, needs no far-end answering equipment, and measures continuously 24 hours a day, every day. This provides information not only about the reliability of the network, but also information useful to network planners (through destination analysis) and information about customer behaviour, which is invaluable when assessing the impact of new services such as Call Waiting and Caller Return.

Choice of sampling rate

The choice of n , the sampling rate, has major implications not only on the statistical reliability of any information derived, but also on the volume of data that is collected and processed. There are typically 200 million call attempts in the UK daily. Sampling 100% of calls would result in the exchange ceasing to switch calls, in favour of becoming a data collection system. Sampling too few calls would result in calls to some destinations rarely being sampled (for example, international payphone calls from Skye to Mauritania). Clearly a compromise is necessary, and after a number of iterations, the standard national sampling rate has been set at 1 in 300. Statistical studies have underwritten this choice as sufficient to provide a robust sample of most call types. Note that the MAC system would not have measured to such a high level of accuracy nor could it cover every possible destination.

Requirements for new system

Once the decision has been taken to consider the replacement of any system, it is necessary to review the functions that are expected of that system, to ensure that it can be made as future proof as possible. One of the key requirements for the new system

was that it should be able to measure the performance of the plethora of new services being introduced. It is clearly unrealistic to develop a new system for every new service, therefore there is a need to develop a generic approach. The system must be flexible enough to meet the needs of a variety of different groups interested broadly in 'performance' but have differing, detailed requirements (for example, product lines, network manager, regulators, customers, etc.). Experience had already demonstrated that a great deal of the required information could be extracted by examining live call information. The extraction of meaningful information, however, is greatly dependent on having an effective information processing system and this is described in the next section.

Processing Live Call Samples

The detailed call information is not in a format that can be used directly and so it needs off-line processing to render it into a usable form. The data is collected and stored at each processor site, and is polled by one of the network mediation processors (NMP) daily. The polled data is then written to magnetic tape and transferred to the computer centre, where it is loaded onto a mainframe computer. This machine is the host for TXDOP (digital exchange output processor), the most important calls database. The complete processing system is shown schematically in Figure 2.

At the tape interface, incoming tapes are processed to extract the data from the different exchange types (System X, AXE10, TXE4) and dispatch it to different destinations, including to the main calls database.

Data is also extracted at the tape interface to supply other systems, including SXTR (a traffic recording system), TXRA (which is used for estimating call revenues), and OTIS (used by capacity planners).

There are four main databases:

- **N3**—is the main database which holds all call records sampled at 1 in 300.
- **NN**—holds call records from sampling rates other than 1 in 300. This is used to identify exchanges which have had the sampling rate incorrectly set, and for special studies.
- **BTI**—holds a subset of 1 in n call records for all outgoing international calls (that is, those calls with dialled digits beginning '010' or '00')
- **RESTARTS**—holds details of automatically logged exchange restarts, data which is used to assess the reliability of the data. This is a record of the date and time of events occurring on exchanges that will have resulted in a temporary suspension of call sampling.

Depending on the authorised user access, any of the databases can be viewed and analyses undertaken. A number of extract programs are run against the main calls database (N3), which has now become an important source of customer behaviour and network performance.

Nodal analysis for digital exchanges (NADE) is the largest extract, running each week and at the end of each calendar month. This extract, which is fed to each exchange performance information collection (EPIC) system, provides the field with visibility of call-failure statistics, and is the functionality that has replaced the MAC-based test call sending measures. Failure rates, by individual processor site, can be combined by EPIC to produce zonal performance in various combinations such as geographic or organisation based. The actual detailed call records of failures due to BT (formerly plant engaged and plant defect) are also included in the extract for transmission to the

EPIC systems, which allows more detailed analysis for trends and patterns.

District call destination analysis extract (DCDA) is used to assist planners by providing an analysis of the dialled digits to determine the destinations of calls.

Call route analysis and factoring toolkit (CRAFT) extract is used to support the production and analysis of network costs for products and services and contributes to negotiations with other licensed operators and OFTEL.

It might be useful to get a feel for the scale of the processing required. Of the 200 million call attempts daily from local exchanges, records from one in every 300 find their way to the main calls database. Additionally, sampling occurs at trunk exchanges and the combined affect is to contribute over 1 million call samples each day. The database has been designed to hold 75 days of data, which, with an allowance for growth, requires capacity for 125 million call records. Along with the necessary reference tables this requires a database storage size of 53 gigabytes. The design and development of a database of this size has not been without problems, some of which are described in the next section.


Implementing the Change

A number of factors needed to be considered when implementing the change. These factors were technical, cultural and organisational.

Technical aspects

The major problems here were connected with the physical design and realisation of the database. The biggest challenge was the sheer scale of the database. BT has the reputation for operating on a megalithic scale; for example, BT has the largest vehicle fleet in Europe, and is one of the biggest spenders of capital investment. The calls database is no exception and is probably the largest DB2 databases of this type in Europe.

Figure 2—Schematic of the processing system



The 'scale factor' has meant breaking new ground in developing and operating the database. The most difficult aspects have been in the performance area where BT has had to make some trade-offs and compromises to optimise the speed of interrogation and extraction of data. This led to a project overrun of twelve months and a reduction in the initial functionality of the system. There have inevitably been some painful learning points. Some of them are listed below and, although they may appear obvious, they can be easily overlooked in a project of this scale.

- ensure that the requirements are clearly defined and understood by customer and supplier at the outset;
- pilot the proposed solution;
- rigorously apply strict project management principles at all stages;
- prioritise the key deliverables; and
- if you have to compromise, phase the implementation, addressing key deliverables first.

Cultural aspects

Network and exchange management within Zones relied on MAC-based information for many years, and a culture had developed which continued to insist that visibility of real-time MAC failures was essential for the efficient running of the network, despite the significant reduction in failures being identified by the early 1990s. There was also an expectation that failure rates should continue to decline as they had in the past, and that network administrations should continue to be active in achieving that reduction. Another major concern was that, without MAC, there would be a loss of early visibility of major network failures, which would have a direct and deleterious effect on customer service.

Organisational changes

As mentioned earlier, an extensive organisation was in place to support the MAC processes. The introduction of live call based measures offered the opportunity to centralise much of the analysis work. BT has now reduced the number of MAC machines from 63 to 21, and with this reduction there has been a consequential reduction in people from almost 400 people to 25. This significant saving of people has been a major success and has helped the field staff achieve their operational targets.

Field trials

To address the very real concerns of a number of operational groups, it was necessary to demonstrate the viability of the new system, prior to obtaining full authority to effect the changes. With the cooperation of management and staff in London and Scotland a three-month trial was held, during which no MAC-based information was provided to the real-time duties within the two network operations units (NOUs). However, before the start of the trial, each Zone worked with the headquarters unit in identifying alternative sources of data that were already deployed within their Zones. This was to ensure that there were no unnecessary blind spots in the ability to detect network problems. During the trial, if a major failure was suspected, an escalation process was brought into play which was designed to establish whether the NOU was aware of the problem without making it known that MAC had seen a suspected major failure. All such events were logged for subsequent analysis. In the event, all major failures were recognised by the NOU before the MAC staff were aware of any problems. The trials were extremely successful and, on that basis, full implementation was authorised.

Major differences in measures

There are a number of fundamental differences which need to be borne in mind when comparing results from

MAC-based information and from live call-based measures.

In the event of a major failure, MAC will 'see' it if it occurs between 0800 and 2100 Monday to Friday. As test calls are sent every 7 minutes during the working day, there is a chance that a short failure might be missed. During normal operation, the level of call failures on a digital network is extremely low. However, in the event of a major failure at a digital exchange it will almost certainly result in a cessation of the measurement sequences (this information will be logged as 'missing data'). However, customers calling that exchange will begin to experience difficulties, resulting in a mushrooming of repeat call attempts, most or all of which will fail. The increased level of call attempts automatically will cause an increase in calls sampled, again most of which will be logged as failures. Thus the actual customer experience will be reflected.

Future Applications

The access to live call sampled data provides an administration with a significant amount of information about its network and the way its customers use it. As the general level of network failures is extremely low, it becomes less cost effective to concentrate on failures of this type. Customer behaviour, however, is an increasingly useful area of study, and data from live call sampling has been used, for example, in connection with the national code change project, to establish how the public has moved from dialling 0- to 01-. It has also been used to chart the changes in the level of 'Ring Tone, No reply' calls during the early period following the launch of *Caller Return / Caller Display*. However, it must be recognised that most people will have some difficulty distilling the information from 125 million records, without a little help. The first level of help is in framing the enquiry into an unambiguous one, recognising that there

are over 1000 different failure reasons. But even when that question is unambiguously framed, the answer still may not be straightforward. A very powerful tool has been demonstrated recently by BT Laboratories, using visualisation techniques⁴, where data is displayed in a number of flexible three-dimensional formats in colour. This approach has helped to identify anomalies in network and customer behaviour and certainly will be used to significant effect in studies of the future.

References

- 1 SAPSFORD, B. G. Measurement and Analysis Centres. *Post Off. Elect. Eng. J.*, Jan. 1978, **70**, p. 243, and April 1978, **71**, p. 28.
- 2 OODAN, A. Quality of Service and Network Performance. *Br. Telecommun. Eng., Structured Information Programme*, Section 13.1, April 1992.
- 3 MULLEE, T. and OODAN, A. Telecommunications Quality of Service: Principles and Management. *Br. Telecommun. Eng.*, April 1994, **13**, p. 53.
- 4 TATTERSALL G. D. and LIMB P. R. Visualisation Techniques for data mining, *BT Technol. J.*, Oct. 1994, **12**, p. 23.

Biographies



David Hand
BT Networks and
Systems

Dave Hand joined the British Post Office as an apprentice in 1965. He is a Chartered Engineer, a Member of the IEE, a graduate of the BT Telecommunications Masters Programme and obtained an M.Sc. in Telecommunications Management at the University of Wales. He has worked in a wide range of jobs in the field of network planning and operations, and he currently heads the Product Launch, Network Performance Management and Private Services Unit in Networks and Systems.



Deryck Rogers
BT Networks and
Systems

Deryck Rogers joined the British Post Office as an apprentice in 1964. He has been closely involved in all aspects of network performance measurement and management for a number of years, firstly on TXE4 and more recently on digital performance. He currently leads the Network Performance Systems Team in Networks and Systems. He obtained a B.A. degree in 1982.

Eliminating the Handicap of Special Needs

People with motor, sensory or mental impairments are part of a broad population who can benefit from telecommunications services. Designing for all potential users—not just the theoretical average—reduces costs, improves ease of use for everyone and ensures that people with special needs are not handicapped by unfriendly technology.

Introduction

Simple features can have a major impact on the accessibility of a telecommunications product or service, for example:

- A pip on the '5' key helps blind people use a keypad.
- Increased spacing between keys enables people with minor dexterity impairments to operate them without accidentally pressing adjacent buttons.
- A notch on the side of prepayment cards enables a blind person to orient it for entry into a card reader.

Designing user-friendly systems needs sensitivity to users' characteristics, limitations and requirements, and the application of *user-centred* methods¹. BT has applied this approach to numerous products and services with features specially tailored to those with special needs.

Special needs refers to the requirements of people with significant motor, sensory or mental impairments, although some groups argue for extending it to include young children and people who are vulnerable from social causes².

The size of the population with special needs in the developed world is large and growing. The United Nations estimates that, on average, 10 per cent of a nation's population have a disability of some sort, although it is hard to estimate the number accurately because many avoid being categorised as disabled while others want to be counted and visible. One study put the number in the UK, with a population of about 56 million people, at 11.6 per cent

or 6.5 million³. This number may be an underestimate as it excludes dyslexics and people with learning difficulties who could have difficulty with information technology designed for the average user. A study of 13 major European countries published in 1990⁴ estimated a total of 50 million disabled people—a group almost the size of the entire population of the UK. The number of people with special needs is increasing with the growth in numbers of old people, since changes related to age are the largest cause of impairments.

We all find it difficult to use a product or service if the user interface fails to take our limitations into account. A motor, sensory, or mental impairment only becomes a handicap if it prevents a product from being used. A system prevents an impairment from becoming a handicap if it does not interact with the impairment or includes features to compensate for it. Such a design can help all users get the best from the product.

People with special needs are simply part of the broad range of human variation. Designers who believe that they are representative of end users, and so design for themselves rather than for all users, can exclude large numbers of potential customers. People vary significantly in many ways, including variations due to impairments. If you are tall you have a different reach than a short person or someone in a wheelchair, and a young engineer may have very different mental skills from an artist, accountant or senior citizen.

Designing for Special Needs

Identifying and meeting users' requirements and implementing



Designing for all users ensures easy access

Deaf people discussing a prototype



design improvements becomes increasingly expensive as a development project progresses. The full range of end users should be considered at the earliest stage. The processes are similar whatever the target user group, though minor modifications may be necessary for working with people who have an impairment.

It's good to talk

There is no substitute for talking to customers, and people with special needs are no exception to this good practice. It is easy to consult disabled people because of support from the government departments and non-government organisations that provide services for them. These organisations will often identify suitable people to take part in requirements capture, usability testing or field trials of new services or products. They may also be able to advise on how to contact people with a specific impairment and may have facilities—such as day centres—where members meet on a regular basis and could try out prototypes or hold group discussions.

Awareness training

It is important to be sensitive to the needs of customers. Organisations such as Action for the Blind run awareness courses about specific disability groups. Awareness training is a worthwhile investment and helps the design team to work well with people and enables them to pick up points raised in evaluation discussions or interviews that they might otherwise miss. For example, basic sign language training can improve the attitudes and sensitivity of a project team to the problems of deaf people as well as helping to communicate with them. Developing a good relationship is vital so that it is possible to try out a potential system more than once.

Conceptual design

The requirements of potential customers should be considered when the concepts of a new system are being developed. The aim is to define

what a product should offer and the usability targets it should meet. Interviewing existing users, analysing ways in which people use similar products and studying how they cope without the proposed system are all useful approaches to gathering information relevant to usage and usability. This also helps the development team identify and interpret design recommendations⁵.

Capturing user requirements provides the foundations for a successful design, though several iterations and the stimulus of early prototypes may be needed. Dealing with special needs is, in some cases, simply a question of including adequate representation of elderly or disabled people within the population considered⁶. For example, terminal hardware selection for an information system should consider the connection of devices such as speech synthesisers for blind operators.

Iterative improvement

Trying out a simulation or prototype system with potential users reveals ways of improving ease of use and allows user performance to be compared with targets. Improvements can be made to the prototype system and supporting documentation followed by another cycle of improvement. The power of personal

Alexander Graham Bell—Like many products in everyday use, the telephone is a 'spin off' from research into technology to assist disabled people. Bell invented the telephone while working on a hearing aid for his wife.



The BT usability laboratory—Flexibility in the layout of a usability laboratory enables controlled experiments to be carried out that make use of realistic simulations of the tasks equipment will be used for. This flexibility should include access for disabled people

computers enables us to take sophisticated prototypes out of the laboratory to wherever it is most convenient for users, such as a day centre for a specific group of disabled people.

Laboratory testing enables usability to be measured scientifically. Potential users carry out representative tasks with the system under controlled conditions and performance and opinion data is gathered. The effects of unwanted

variables are eliminated by careful planning and statistical techniques. In addition, valuable insights into improvements come from people's informal comments about the system. It is easy to include people with disabilities in these evaluations although minor changes of procedure are sometimes necessary. People with no vision need spoken or Braille instructions or questionnaires. People who are deaf from birth may not have

Claudius II—People who can hear but cannot speak have great difficulty using the telephone. BT has developed a product that will generate short spoken messages at the press of few buttons to help these people. The design team carefully considered the end users particularly in relation to the controls, layout, panel marking and operating procedure.





A textphone—To communicate over the telephone network, some deaf people use a telephone with a keyboard and display to transmit and receive text messages. When setting up a call with a textphone, a deaf customer cannot hear the spoken information messages from the network. Investigations were carried out to identify the best structure and content for text messages that could be provided to textphone users in place of the spoken word.

well developed written English skills and so simple language supplemented with good graphics should be used. Laboratories should be 'wheelchair-friendly' and have arrangements to alert deaf people if there is a fire alarm.

Laboratory trials take place in a rather artificial situation. Field trials forego the rigorous control of the laboratory, but have the advantage of exposing the system to representative users in much more natural circumstances. This is a useful way to find out how a system will fit into the lives of real customers. Self-reporting diaries, discussion groups of users and most other ways

of monitoring a trial will be applicable to participants with a disability. It may be necessary to tailor support or monitoring arrangements to the abilities of the participants. Perhaps the most important aspect of setting up a field trial is to ensure that the system has real value to the participants so that they wish to use it. With communications services this means selecting participants who really want to communicate with each other.


Benchmarking

Comparisons between products make it possible to identify which are

strongest and where there are opportunities for improvement. Checklists of important system parameters make it easy to compare products⁷. Wherever possible, the comparison should include evaluation and consultation with the disabled people who will be using the product. Checking against recommendations that are based on current practice is not always adequate owing to the rate of change of the technology.

Concluding Comments

It is more than seventy years since telephone engineers first modified switchboards to enable disabled ex-servicemen to operate them. Providing access to products and services by people with special needs is an important aspect of improving service for all customers. The rapid advance of technology promises many exciting opportunities—the plain old telephone is evolving rapidly into a multimedia terminal. Such developments present opportunities and threats for people with special needs. Increased availability of visual communications allows deaf people who sign to communicate more effectively, but emphasis on visual display of information may cause difficulties for those who are blind. The complexity of new services may make them difficult to master for those with learning difficulties or the



(Left) Wearable computing and communications system—Some blind and elderly people find travelling even short distances difficult, becoming disoriented or lost even when close to their destination. A wearable computer may be the answer to these travel problems. It could determine where it is from positioning technology, communicate with databases or service centres and guide the user by means of synthetic speech. BT is working on wearable computing and communications systems and participating in a project on applications for blind people within the European Technology Initiative for the Disabled and Elderly.

educationally disadvantaged. On the other hand, multimedia systems will support many of the activities of daily life through new services such as tele-shopping, tele-medicine and tele-education. These systems offer the potential benefit of making it easier to interact with others and information sources, effectively removing the disadvantages of speech problems, visual disabilities and other impairments^{8,9}. Dealing with the human factors of these new systems presents an exciting challenge and opportunity to extend our society in the widest sense.

Acknowledgement

The author would like to thank the many people working to eliminate the handicap of special needs—particularly those with disabilities who have freely given up their time to participate in discussion groups and equipment evaluations.

References

- 1 COOPER, M. B. Human Factors in Telecommunications. *Br. Telecommun. Eng.*, July, 1994, **13**, p. 98.
- 2 The Impact of Developing Technology on Vulnerable Groups. Final report of a BT Issues Exchange, Nov. 1994. BT Corporate Relations, Procter House, 103 West Block, 100-110 High Holborn, London, WC1V 6CD.
- 3 FREDERIKSON, J. Statistics, In Stephen Von Tetzchner (ed.) *Issues in Telecommunications and Disability*. Office for Official Publications of the European Communities, 1991.
- 4 SANDU, J. S. and WOOD, T. Demography and market sector analysis of people with special needs in thirteen European countries: a report on telecommunication usability issues, 1990. Special Needs Research Unit, University

of Northumbria, Newcastle upon Tyne, NE7 7TW, UK.

- 5 ETSI Human Factors (HF) Committee. Access to telecommunications for people with special needs—Recommendations for improving and adapting telecommunications terminals and services for people with impairments. Oct. 1991, Technical Report ETR 029, European Telecommunications Standards Institute, B.P.152, F-06561 Valbonne Cedex, France.
- 6 FURNER, S. The user with special requirements in the application of structured design, *Proc. Access to Computer Systems by Blind Persons (including Windows, Icons, Mice and Pull-down Menus)*, Madrid, Spain, 31 May–2 June 1989.
- 7 GARNEY, K. Choosing the right telephone. *Quarterly Newsletter of Danish Centre for Technical Aids for Rehabilitation and Education*, No. 2, June 1992. Danish Technical Centre, Graham Bells Vej 1 A, DK-8200 Aarhus N.
- 8 SMYTHE, P. Uses of Networked VR to Assist People With Special Needs. *Proc. Virtual Reality and Disability*, Oct. 1993. Swedish Institute for Computer Science and Department of Education at Stockholm University.
- 9 MILLER, S. K. Video Goggles Give Second Sight to Ageing Eyes. *New Scientist*, 23 May 1992, No. 1822.

Bibliography

TETZCHNER, S. *Issues in Telecommunications and Disability*, 1991. Office for Official Publications of the European Communities L-2985 Luxembourg.

LAZZARO, J. J. *Adaptive Technologies for Learning and Work Environments*, 1993, ALA Books, 50 East Huron Street, Chicago, IL 60611, USA.

Biographies



Stephen Furner
BT Networks and Systems

Stephen Furner joined BT as a trainee technician working on the development of the public Telex network. He contributed to the development of time-division multiplexers for Telex transmission and the replacement of elements of electro-mechanical Telex exchanges with digital technology. He left BT on special leave in 1975 to study for a degree in psychology at the University of Keele. On his return, he joined the Human Factors Unit where he works as an engineering psychologist. He has led research and development projects, carried out consultancy, represented BT at international standards bodies and works on numerous European research programmes. He has published over 20 academic papers, reviews contributions to major journals and conferences, and is an Incorporated Engineer and a Chartered Psychologist.



Martin Cooper
BT Networks and Systems

Martin Cooper joined the Post Office in 1967. In 1971 he graduated from Southampton University with a B.Sc. in Electronic Engineering and gained an M.Sc. in Work Design and Ergonomics in 1975. He is a fellow of the IEE and the Ergonomics Society. Since graduation, he has worked on a variety of projects in the human factors field. In 1988, he was awarded the Sir Frederic Bartlett medal of the Ergonomics Society for his contribution to the application of ergonomics in industry. His current role is managing a team of psychologists, ergonomists, designers and engineers studying ways of matching future systems to the needs of their users.

Challenges in Information Visualisation

We are surrounded by an ever-growing and challenging world of data. However, the value of this data is not intrinsic, but lies in enabling us to make better informed decisions and in increasing our shared knowledge and understanding. This article explores the critical role for visualisation in bridging the gap between the abstract, analytical world of data and the digital computer, and the real, analogue world of human problems and experience. The discussion is structured around four challenges in information visualisation.

Prologue

The revolution in information systems is presenting technological challenges in areas such as data storage, transmission and security, but perhaps the greatest challenge is in retaining the intimate involvement and understanding of the end user.

While computers are becoming ever more powerful, they are only an aid to problem-solving, and there remains a critical role for intangible human capabilities such as intuition, experience and imagination. This article considers the potential for visualisation to bridge the gap between the abstract, analytical world of the digital computer, and the human world. The traditional interface of mouse, keyboard and screens of text allows us to work **on** computers, while techniques such as visualisation will truly enable us to work **with** computers.

The overview reinforces some of the trends in data and business that are shaping the 'Information Age'. These trends are drawn together into a framework mapping the conversion of data into decisions and knowledge, and the framework is used to structure the remainder of the paper around four major *challenges* in information visualisation: a growing volume of data with declining information content; more complex data analysis tools and models; increasingly abstract data and issues; and a wider, less specialist audience.

Overview

The *Information Age* is a widely used term for the current era in which we are overwhelmed by data and information in diverse media and

'the single biggest problem we face is that of visualisation'

Richard P. Feynman
(Los Alamos, 1945)

formats. Paper records, journals and books are increasingly supplemented, or even supplanted, by computer-based storage and retrieval, with vastly increased capacity, accessibility and variety of content. Moreover, the use of computers in monitoring, modelling and simulating enables problem-specific data to be collected or generated often at minimal cost. In the Information Age, lack of data is rarely a problem.

In contrast, it is increasingly common to encounter problems for which all the necessary data is potentially available to make a complete assessment of the risks, and hence to reach an optimum solution. However, mapping that data through to timely, quality decisions has never been harder. Figure 1 shows a possible framework for the conversion process, and identifies some of the challenges.

The first challenge is locating and retrieving the relevant data, from a data world which is growing in size, but is also declining in average information content, as the provision of data becomes ever easier and cheaper. Analysing and interpreting data is the second challenge, and although sophisticated computer-based tools and techniques are available, they can often appear as inaccessible 'black boxes' to the uninitiated. Even when the appropriate information has been extracted, two further challenges remain: the final conversion through to real




Figure 1—Conversion of data to informed decisions and shared knowledge

benefits, expressed in our framework as better informed decisions and the sharing of increased knowledge. Difficulties in these areas centre on the intangible nature of many of the issues and information under consideration, and hence the problem of establishing a common and comprehensive perspective. The four challenges are closely interrelated, and in any specific scenario the boundaries will be blurred if not invisible. However, the framework is nevertheless useful in clarifying some important issues and potential solutions.

The *challenges* are driven by the increasing power and accessibility of computing and telecommunications, trends which are certain to continue and even accelerate for the foreseeable future. The same advances that are giving rise to the challenges will also provide partial solutions, through developments in data mining and information agents, for example^{1,2,3}. However, visualisation has an essential role to play by providing new and powerful tools to help us understand and solve complex, data-intensive problems. The remainder of this article explores and illustrates that critical contribution.

Exploring, Navigating and Browsing

'where is the knowledge that we have lost in information?'

T. S. Eliot

Data is an increasingly important strategic corporate resource, but, ironically, advances in collection, storage, transmission and duplication

facilities have in many respects also devalued data. Just as the advent of the photocopier, the word processor and the desktop printer led to a proliferation of paper without a corresponding step change in the information content of the average in-tray, so the ease with which electronic data can be generated has produced a comparable decline in quality.

A decade ago the provision of electronic data was a highly specialised business run by professionals offering a quality product and charging premium prices. Today, there are countless new and unproven information providers, and it is almost as easy to supply data as it is to receive it⁴. Electronic news groups and bulletin boards may contain up-to-the-minute details on the latest advances in a specific field, but they are invariably swamped in mundane trivia, and such systems are no substitute for refereed papers in an established journal. There are significant potential costs associated with poor quality or inappropriate data, and it is increasingly a case of *caveat emptor*, let the buyer or receiver beware.

Over the same decade, our ability to consume information is largely unchanged. The value of this growing data world will therefore only be realised if it is accompanied by developments in tools to help us locate and retrieve **appropriate data**. Data will be available to minimise the risk in most decisions, but will it be cost-effective or even feasible to retrieve that data? The *infocruiser* of the future will be required to navigate efficiently to retrieve specific data, to explore data

surrounding a topic of general interest, and effortlessly to browse the wider data world.

This *infocruiser* will be equipped with sophisticated computer-controlled navigational systems which will adapt both to our personal profile and to the task in hand. There will be extensive use of pre-computed catalogues and resource maps; automatic filtering and summarisation will reduce the complexity of the data landscape; and increased telecommunications bandwidths will raise the speed limit. However, retaining the full involvement of the human driver, working in concert with the automatic systems, will greatly increase the flexibility of the craft. Visualisation will present the data landscape in a natural and intuitive form, making full use of the human capacity to absorb and interact with complex images, far beyond typical directory structures or screens of text and numbers.

A concept demonstrator has been developed to illustrate the use of an interactive visual interface for exploring network fault data. The visualisation (Figure 2) allows an experienced network manager to browse the database looking for trends or predictive patterns. Geographic and temporal views are readily interchanged, with a range of flexible filters to help pinpoint specific problems or 'blackspots'. Preset thresholds and filters can be used to identify standard exception conditions—the autopilot on our *infocruiser*. The interactive visual interface makes effective use of the knowledge and experience accumulated by the network managers, and retains their **involvement** in the process.

It is only a small step to extend visualisation to an immersive experience in which we make full use of the human senses to acquire, process and interpret the data world⁵. Other work makes use of emotional icons to increase the effectiveness and impact of the computer interface onto a complex data landscape (Figure 3). The icons respond to the presence of a




Figure 2—Network alarm data

user in this virtual world with characteristic behaviours, which might depend on the profile of the user and their current activities and interests. For example, icons will advance or retreat, grow or shrink, act aggressively or passively, reflect-

ing the relative importance and attributes of the data or action that they represent. The visual impact is enhanced by the use of sound to create a rich information environment.

We therefore conclude that in response to the first *challenge* of

growing data volume combined with declining information content, there is a key role for visualisation in enabling human exploration, navigation and browsing of complex information environments. A natural and intuitive visual interface can retain the critical contribution from human perceptual skills and ensure that opportunities for lateral thinking, or perhaps an unexpected leap of imagination, are not lost. Programming a computer to 'look for something interesting' in a database is a major undertaking, but given the appropriate tools, it is a task for which humans are well equipped.

Understanding Complexity

'the purpose of computing is insight not numbers'


R. W. Hamming

The second *challenge* is the conversion of *appropriate data* to *relevant information*. There are many parallels with the initial extraction of the data in that the overall goal is summarisation and compression without significant loss of content. However, in this step we also include more sophisticated processing tools which aim to enhance the value of data through analysis and interpretation.

Complex data modelling and analysis tools are increasingly available, but the sheer sophistication of these 'black box' solutions can often be their undoing. The range of options and facilities can be daunting for the novice, and only a trained and experienced user can obtain the optimum performance. Moreover, response to the output ranges from blind trust and acceptance, because of the 'expert' status of the analysis package, through to sceptical distrust, because the underlying techniques are invisible and poorly understood.

Once again, there is scope for visualisation to retain the human involvement and understanding. Consider the simple example of

Figure 3—Emotional icons






Figure 4—Subjective assessment of a telephone connection

regression analysis to fit a straight line. How many of us can claim familiarity with the interpretation of the regression coefficient, the standard deviations and other elements of numerical output from a computer calculation? And yet seeing the points and regression line plotted on a graph provides an immediate understanding of the process. We are likely to

spot weaknesses in the interpretation. Was it appropriate to fit a straight line? Is there unexpected structure in the distribution of the residual errors? Roughly, how good is the fit? Ultimately, we cannot afford to neglect a more rigorous analysis, but visualisation retains our involvement and understanding, and ensures that we do not generate


solutions which are divorced from the original problem.

Practical business models are typically much more complex than the simple example of linear regression, so the need for user involvement is even more apparent. As with the first challenge, the solution lies in combining sophisticated computer techniques with existing and highly developed human skills. Artificial intelligence can perhaps help to guide us through the complexity, by suggesting appropriate models or techniques, but retaining user participation through computational steering would enhance the overall effectiveness of data analysis. Indeed, there is a pressing need for 'intelligent visualisation' which will help novice users select an appropriate representation for their data from a bewildering array of options.

Work on interactive modelling has included complex business models, and a visual interface to a number of models which assess customer perception of telephone connections⁶ (Figure 4). The latter have a dozen or more input parameters specifying the characteristics of the circuit, and the visualisation allows the user to explore perceived performance as a function of selected pairs of parameters. Although the visualisation adds nothing to the capability of the underlying models, the increased user involvement offers the potential for faster understanding and more rapid, exhaustive exploration of the model's predictions and limitations (Figure 5). It is also possible to include experimental data from controlled conversation and listening tests in the visualisation, which enables intuitive comparison between experiment and theory.

As in the first challenge, there are even greater benefits to be realised by exploiting the other human senses⁷, and here there is a natural extension to simulating the telephone connection under consideration. This could be done with a network emulator⁸, which would mimic the circuit conditions at a selected point on the modelling surface. Users of the model

Figure 5—Interactive modelling interface



The successful development and implementation of information systems is perhaps one of the greatest issues currently facing many businesses.

could then experience the speech degradations, and hence develop an even greater understanding of the modelling domain.

One important aspect of the telephone connection visualisation is the combination of different data sources into a common representation. Instead of comparing numbers or graphs in their disparate native formats, using a common visual framework allows much easier identification of discrepancies and relationships. A second example of this is a visualisation of lightning strikes and network alarms, which provides immediate insight into the patterns and correlation (Figure 6)⁹.

A final illustration of the potential for visualisation to enhance understanding of complexity is the representation of mathematical algorithms¹⁰. An example of this is the visualisation of a neural network algorithm which has been developed to assign radio channels to base stations in a mobile communications network¹¹. The usual method of expressing this type of algorithm is by means of a large number of non-linear coupled differential equations, the behaviour of which is very difficult to predict. Moreover, the efficiency of these algorithms is highly dependent on the choice of the free parameters. Through visualisation, the output from the algorithm can be viewed at each computational step (Figure 7) and the user can develop a 'feel' for the internal operation. Visualisation techniques also allow interaction with the model during the development phase, which leads to an intuitive understanding of sensitivity to empirically determined parameters.

These examples illustrate the role for visualisation in the second *challenge* of extracting information from data. Once again, that contribution is achieved by maintaining a high level of human involvement and participation in the processing and analysis of data, thereby retaining the benefit of attributes such as experience and intuition that are




Figure 6—Visual correlation of lightning strikes and network alarms

hard to capture in even the most sophisticated computer models.

Managing Intangibles

'a picture is worth a thousand words'

Having extracted, summarised and analysed our data there remains a further major hurdle in our framework; the data itself has no value, and it is only in the final conversion to improved decisions and increased knowledge that we achieve real benefit. Even if there are circumstances under which human participation makes a minimal contribution to the overall quality of data collection and analysis, that user involvement is nevertheless critical in these final stages. Two distinct *challenges* are identified in Figure 1, and this section considers the problems in managing increasingly abstract problems against ever-decreasing timescales.

The successful development and implementation of information systems is perhaps one of the greatest issues currently facing many businesses. It is an area littered with infamous failures, and despite many years of experience and the introduction of sophisticated management methods, there is no consensus that the performance of the industry is improving. It is common to use analogies with construction or

Figure 7—Evolution of the channel assignment algorithm





Figure 8—Evolution of software metrics

manufacturing to describe the information systems development process with, for example, parallels drawn between the roles of the systems analyst and the project manager; and the architect and the foreman on a building site.

In the implementation phases of both a large civil engineering project and a major information system development, success hinges on the coordination of a multi-disciplinary team around a complex design specification. However, a fundamental difference in the construction project is that every member of the team is able to visualise the physical structure that they are building. Some aspects of an information system are tangible, but a major component is the software which is an abstract and invisible collation of computer instructions. The intangible nature of the product raises a wide range of management issues, from the performance of individual workers through to overall coordination and control. For example, an individual on a software development is less likely to have a broad appreciation of their role in the overall development; they cannot necessarily see the wall that their bricks are a part of, or the building that will be formed from those walls. They are also less likely to understand the overall user requirements, and the way in which they have been broken down to define their individual contribution. The end result is a greatly reduced understanding and involvement in the overall project, and hence a greater likelihood that effort will be misdirected. For the manager, it is also much harder to obtain early visibility of potential problems and delays, and there may be little physical evidence of progress against which to judge short-term success.

As with the problem of understanding complexity in data, a partial solution can be found in more sophisticated computer-aided development tools and management techniques, but many of the difficulties are directly attributable to the



intangible nature of the product. A simple visualisation has the potential to improve, dramatically, control and understanding.

There has been much work on the application of visualisation to software development and management¹²⁻¹⁴, with many commercial packages now benefiting from interactive visual interfaces. However, there remains considerable scope for further advances, and even if a software system is viewed merely as a collection of statistics, many existing visualisation techniques can be readily applied. The display of software code metrics (Figure 8) shows the metrics associated with a software system across five major releases. The most recent is shown on the top plane, with each point representing a module in the system, and each peak representing a metric associated with that module. As the system evolves, it is possible to see changes in the profile. The designer is able to pick out long-term trends; for example, the two emerging spikes could mean the underlying modules are becoming overly convoluted and are in need of re-engineering and simplification.

In a more revolutionary approach, virtual reality is used to create a richer information environment, with shape, behaviour and colour representing attributes of the software (Figure 9). Each module of code is mapped to a sphere floating in free space. The radius of the sphere is determined by the lines of code in the module, and intermodule function calls are shown as cross-connecting links. Change requests are mapped to the rate of spin; the greater the number of change requests, the faster a module spins.

An area closely related to developments in information systems, is the implementation and management of major change in business processes. Again, much of the end result is largely intangible and many problems can arise through an inability to visualise the full impact and complex interrelationships of the proposed changes. Figure 10 shows a visualisation of a process model within BT. The user can quickly obtain an overview by moving freely around the three-dimensional representation, or can track specific components of a process in more detail.




Figure 9—Virtual reality information environment




Figure 10—Process model visualisation

Figure 11 is a final example of the use of visualisation to improve the management of an intangible problem. The visualisation shows the performance of a call centre in handling telephone enquiries, highlighting times when additional staff might be required, or illustrating the impact of an advertising campaign. Where the queues and enquiries are physical—at a post office or bank, for example—efficient management is much easier, in that all concerned can clearly see the situation and will modify their behaviour accordingly. The queues and lost calls on a telephone network are invisible, but in business terms they are equally important.

The key role for visualisation in this third *challenge* has been in retaining an overview of an abstract problem, and in providing immediate visibility of changes and developments which might not otherwise have an impact for weeks, months or even years. Human skills, such as peripheral awareness, can greatly improve the quality of management and decision making at all levels, but this requires the problem and associated data to impinge directly on our senses.

Communicating a Vision

'I hear and I forget, I see and I remember, I do and I understand'
Old Chinese proverb

The final *challenge* is in effectively communicating a vision such that it is accurately shared and understood by the audience. This seemingly easy

task can employ an entire marketing department in defining and conveying the less tangible attributes and benefits of an overall product offering. When we are faced with conveying the features of a service, the use of examples or even role-playing will typically be used to create a more substantive reality. However, we are increasingly required to manage and communicate processes, and to design and develop products and services

which have no meaningful physical embodiment.

The development of information systems can again provide an illustration of the difficulties in establishing a common view of a non-physical product. Many problems can arise in developing the initial specification which maps the required business benefits through to specific functionality in the final implementation. It is unlikely that

Figure 11—Calling patterns






Figure 12—Three-dimensional network visualisation

business managers and IT specialists will be able to find a common language for their discussions, and it is possible that they will develop widely differing views and expectations of the same system. They are unlikely to share engineering drawings or a computer model of the final system which might enforce a common understanding, albeit from a different business perspective.

Similar considerations apply in agreeing the specification for a telecommunications service, such as designing a complex internal company network. The technical expertise of the network planner and the business needs of the customer can find common expression in a three-dimensional visual representation (Figure 12). This visualisation can be much more than a colourful interface, with a range of feedback and interaction points to encourage active customer involvement and hence enhanced understanding. The aim is to demystify the design process by providing an intuitive interface, while retaining the necessary underlying rigour.

Moving on from establishing a common view on a specific product or service, another area in which problems often arise is communicating a long-term vision or strategy. There is often a wide audience that has not been actively involved in background debate or discussion, and who tend to be disinterested, or at best sceptical. An effective visualisation not only provides a common language in which to express the vision, but can signifi-

cantly enhance the impact that can be achieved by spoken or printed words alone. Concepts and new ideas that can be portrayed visually will reach a much wider audience, and are more likely to be further discussed and considered.

An appropriate virtual world can be particularly effective at conveying the impact of future systems. Commercial virtual world authoring packages, and the power of desktop computers, make the construction of such scenarios a relatively straightforward undertaking. Figure 13 is a screen image from the *intuitive call centre manager*, a virtual world built around a vision of the future in network planning and management.

We can conclude that there is a clear role for visualisation in communicating increasingly abstract visions and product offerings to an audience which would otherwise struggle to find a common language for discussion and debate. An effective visualisation can provide a shared, but perhaps also customised, view of the same underlying data, and has unparalleled potential to make a memorable and lasting impact.

Discussion

At the outset of this article a framework was proposed for the conversion of raw data through to real benefits, expressed as improved decisions and increased knowledge. The steps in that conversion process have been examined, and, using a wide range of examples, some of the main trends



Figure 13—The intuitive call centre manager

and developments have been identified. The overriding theme in these discussions has been the importance of retaining human involvement in all aspects of data analysis and interpretation. Through computers, we are both creating an ever-growing world of abstract data, and developing the sophisticated tools and techniques that we require to interact with it. However, there is a tendency for these 'black box' solutions to neglect potentially critical human capabilities, such as intuition, peripheral vision, spatial awareness and experience. These are skills that we use naturally and effectively in our interactions with the complexities of everyday life, but which are usually excluded from computer worlds by limitations of current interfaces. Visualisation could enable our active involvement and participation in these worlds, and greatly increase the value ultimately derived from the data.

Turning finally to the involvement of network and service providers: what are the implications of this increasingly abstract and data driven world? Once again, the key lies in the importance of retaining human participation in the process of analysis and decision making. Computers routinely communicate with each other at many megabytes per second, but effective human involvement in such transactions hinges on interfaces which use techniques such as visualisation. Visualisation and related immersive technologies will therefore be critical enablers for many future high-

Visualisation and related immersive technologies will be critical enablers for many future high-bandwidth services involving human-computer and human-human communication.

bandwidth services involving human-computer and human-human communication. It will be important to understand, and even to anticipate, the bandwidth and management needs of emerging service providers. Key issues will include the appropriate distribution of resources, such as data and processing, between the service provider, network infrastructure and customer terminal. A range of trials aimed at extending the effectiveness of remote visualisations and shared virtual worlds from specialist physical science applications to broader commercial and domestic markets will be required.

Conclusions

By examining the conversion of data through to improved decisions and increased knowledge, we have provided a glimpse of the potential for visualisation to increase the bandwidth between ourselves and an ever-growing and ever-changing world of data. In this exploding world of abstract data, there is great potential for visualisation to combine long-established human strengths and capabilities with the overwhelming processing power of modern computers. Ultimately, we can envisage a 'mind amplifier' which augments our natural analogue (intuitive and heuristic) skills with the precision, speed and certainty of a digital (logical but unthinking) counterpart.

Before this vision can be realised, we need a great deal more experience in issues such as representation of data, the automated creation of effective information landscapes, and appropriate interaction mechanisms. Real-world analogies such as the *infocruiser* are a valid starting point for these experiments, but we should not be constrained in our exploration of new metaphors and completely artificial environments. Creating a world in which humans and computers are equally at ease in the analysis and interpretation of data will not be easy, but the rewards will be immense.

Acknowledgements

The author would like to thank colleagues at BT Laboratories for their contributions to the work, and for detailed comments on the article.

References

- 1 FRAWLEY, W. J. *et al.* Knowledge discovery in databases: an overview. *AI Magazine*, Fall 1992.
- 2 RICH, ELAINE. Artificial Intelligence. McGraw Hill, 1983.
- 3 FLAVIN, P. G. (ed). *BT Technology Journal* theme issue on Computer Aided Decision Support, Oct. 1994, 12(4).
- 4 BARNES, J., CHALMERS, J., and PEARSON, I. The European Challenge to Telecommunications in East and West, *Proc. 33rd FITCE Conference*, 1994, pp. 72-79.
- 5 KALAWSKY, ROY S. The science of virtual reality and virtual environments. Addison-Wesley, 1993.
- 6 SOUTH, C. R. In-service, non-intrusive assessment of telephony channels—an introduction. *BT Technology Journal*, April 1992, 10(2), p. 13.
- 7 COULL, T., and ROTHMAN, P. Virtual reality for decision support systems. *AI Expert*, Aug. 1993, pp. 22-25.
- 8 GUARD, D. R., and GOETZ, I. The DSP network emulator for subjective assessment. *BT Technology Journal*, Jan. 1992, 10(1), p. 162.
- 9 WALKER, G. R. *et al.* Visualisation of telecommunications network data. *BT Technology Journal*, Oct. 1993, 11(4), p. 54.
- 10 TATTERSALL, G. D., CHICHLAWSKI, K., LIMB, R., and TOTTON, K. A. E. Feature extraction and visualisation of decision support data. *BT Technology Journal*, July 1992, 10(3), p. 110.
- 11 LOCHTIE, G. D. Frequency channel assignment using artificial neural networks. IEE International Conference (ICAP '93), April 1993.
- 12 MYERS, B. Taxonomies of visualising programming and program visualisation. *Journal of Visual Languages and Computing*, 1990, 1(1).
- 13 ROMAN, G. C., and COX, K. C. Program visualization: the art of mapping programs to pictures. International Conference on Software Engineering, no. 14, 1992.
- 14 COZENS, M. *et al.* Visualizing and querying software structures. International Conference on Software Engineering, no. 14, 1992.

Biography



Graham Walker
BT Networks and
Systems

Graham Walker joined BT as a sponsored student and after graduation from Oxford University in 1986 spent six years researching into coherent optical transmission systems. He was involved in the world's first field trial, and made advances in aspects of polarisation control and the noise performance of optical amplifiers. This work resulted in numerous publications, and the award of a Ph.D. from Cambridge University in 1992. More recently, he has been leading the Information Visualisation Group within Advanced Applications and Technologies, exploring novel applications of visualisation to telecommunications data. Current interests include the visualisation of call record data, and of scenarios in business strategy. He is a Member of the IEE.

Ian Pearson

When I'm 64

Technology and Telecommunications Beyond 2025

Computers will gradually become proficient at a wider range of activities previously requiring human intelligence, including their own design. This, coupled with currently expected progress in nervous system connection, may accomplish a direct machine-human brain link by 2025. Enormously powerful machines, using direct brain links to join together the minds of a society, will produce a new infrastructure or Medium. This article explores some of the fascinating possibilities such technology would bring. Although such predictions are highly speculative, exploring the consequences makes it possible to postulate upper bounds on the demands for future telecommunications.

Prologue

What will the world be like in 2025? No one knows. Who in 1964 would have predicted the ubiquity of (or our enormous dependence on) personal computers by today? Predictions over 30 years ahead look like science fiction today, and could easily be off the mark by an order of magnitude or more. Even if a prediction comes true, time-scales may be out by several years either way. Nevertheless, it is necessary to explore and contemplate the possibilities if telecommunications is to get ahead of, or even keep up with, the demands of a society increasingly equipped, aware of, and dependent on, information technology.

Progress in Computing

While today's computers are easily identifiable discrete entities, in the future processing capability and intelligence will be ubiquitous. The vast majority of processing devices will be transparently integrated and unidentifiable as individual machines. However, there will still be a place for machines dedicated to providing computing and communications capability, but these will be very different, and far more powerful than those of today.

Computers have long since overtaken human digital processing capabilities, but we do not yet consider them intelligent (whatever that means!). Indeed, today's supercomputers are thought to have the intelligence of a chicken. However, many of the functions that we associate with human intelligence will be capable of machine implementation by 2010¹. Examples are the use of natural language, equal or better senses including vision and recogni-

tion, understanding text and graphics, and machine use of human memorising, creativity, recognising and learning. Already, computers can beat most grand masters at chess. Some computer architectures and mechanisms will be based on the human brain with a full range of high-fidelity human-like senses, and pre-processing in biosensors where appropriate. In many ways they will be smarter than humans with their interface to the real and computer-generated worlds centred on these new and enhanced senses. By about 2025 it may even be possible to link directly to the human peripheral nervous system through a range of new external and implanted devices¹.

Even though we are currently unable to understand or mimic the working of the human brain, it is probable that computers will be able to emulate, or even surpass, a selection of those tasks associated with human intelligence by 2010 or earlier. This is one of the two key assumptions of this article. The second is the ability to interface directly with the nervous system possibly with a direct brain link, albeit on a longer time-scale.

Just calculators

A few years ago, computers were considered nothing more than advanced programmable calculators. However, by exploiting neural network principles, some are now capable of tasks previously thought to require 'intelligence'. With neural networks (or other technologies), in principle it is not necessary to know the algorithm used by a processor in order to instruct it to solve a problem or make a decision. Our 'cleverer than human' machine would be able to 'understand' the problem with the same level of guidance needed for a

By 2030, computers could have advanced to the stage where a personal computer exceeds most of the capabilities of the human brain

human, and would 'think for itself' as far as solving it. Some authors¹ suggest a 2005 time frame for this being achieved, 2010 is a more conservative estimate.

Technological incest

We can expect these machines to have a significant impact. For example, since it is conceivable that they could design computers better than human counterparts, the rate of development of computing would accelerate, each generation designing even better computers at an even faster rate, perhaps eventually inventing concepts that would be incomprehensible to humans. A typical example of recent human displacement is in printed circuit board design. A human operator may take days to produce and verify a design that a desktop computer can do in hours. As speeds and artificial intelligence techniques improve, machines will take over more and more human tasks such as design.

By 2010, today's electronically reconfigurable devices such as field programmable gate arrays will be very much more powerful and sophisticated. It is conceivable that a computer incorporating these would be able to reconfigure itself dynamically to achieve optimum performance under all circumstances. *In other words, it would evolve to the limits of its hardware technology.*

Computer power exceeds mankind

By 2030, computers could have advanced to the stage where a personal computer exceeds most of the capabilities of the human brain (see Figures 1 and 2). This will be achieved through advances in current technologies such as neural networks, superconductivity, optical, quantum and biological computing and various other fields. Major contributions will also come from technologies which are currently embryonic or poorly understood, for example, nano- or even atomic-scale devices offering full interconnection and high-speed processing.

Digital processor performance

At present, digital computing speed is increasing by a factor of 100 per decade. There are no reasons to see any near-term decline in this rate and some reasons to expect it to increase. Indeed, we could argue that the improved design capabilities offered by each successive generation of computers would give an additional improvement factor of 10 per decade over that resulting from human research alone. Thus, the overall improvement rate could be 1000 times per decade. If computers with the equivalent intelligence of

the human brain are realised in 2010, this means that the 2030 super-computer would be 1 million times more powerful. In addition, let us not forget that major performance enhancements invariably lie dormant owing to architecture and algorithm limitations. When these are liberated, the improvement rate could reach 10 000 per decade. The 2030 super-computer may thus have a power of between 1 million and 100 million times that of the human brain (quite an error margin!). By the same line of reasoning, the PC of 2030 may be equivalent to 10 000 human brains.

Figure 1—Projection of super-computer speed





Figure 2—Projection of memory 'chip' capacity



It is important to note that these figures relate to the speed in accomplishing tasks considered today to require intelligence in terms of memory and knowledge. There remain many areas unique to humans that these figures do not (and cannot—yet!) take account of.

Analogue

Much of the work carried out by 2030 computers may involve analogue processing and computation. The digital approach can of course be used (and in many situations will remain the best option) but the high computational speed and parallelism required for some tasks may only be achievable using advanced analogue techniques. In any case, since our perceptions and interactions with the world are entirely analogue, it seems reasonable to expect that analogue techniques will be used to interface with it and to deal with the problems arising from it. Indeed, many such problems may only be soluble using analogue techniques.

Impact of Powerful Computers

The impact on mankind of the future powerful computers will be staggering. They will have the inherent ability to make better decisions than humans, having instant access to all current knowledge via the world's databases, and having enough storage and processing capability to use it effectively. Many tasks currently beyond our reach could be undertaken with ease. Human roles would therefore change dramatically. The following are a few of the possible outcomes:

- As control of many of the functions of society and business falls to intelligent computers, accounting and finance would become radically different from today.
- Companies would have a more chaotic and amorphous structure, many existing only in the infosphere, with no physical or centralised base.

- Companies will be transformed, reorganised, merged or wiped out to the extent that none of today's companies will still exist in recognisable form.
- The stock market will no longer exist as only equivalent super-computers will be able to compete, and negotiate prices directly.
- It is possible that the whole economy would tend to stabilise under automatic control and co-operation at the highest levels. However, the speed of these computers would allow financial decisions to be made in a time which is short compared to network transmission delays, which could cause stability problems. Alternatively, the 'best' computer in an information region could quickly gain overall control giving regions of local stability.
- With the vast power of these computers, knowledge would multiply rapidly, advancing every branch of science, medicine, and economics, always assuming that we can persuade computers that these are of interest and worth tackling. Where there were gaps in knowledge, computers could initiate new research.
- Computers could be made to be much 'friendlier' than man. On the other hand, the sheer power of these computers will leave mankind intellectually dwarfed, and there will be vast pools of knowledge that only they can understand and use!

Many of the above consequences assert that computers will be capable of making better decisions than man. Here again, the quality of a decision is often in the eye of the observer rather than some absolute property. It may be that there exist some problems in which humans can make 'better' decisions than machines. However, it is equally possible that

appropriate emulation will allow these machines to mimic human decisions in most instances if required. It is certainly true that there are problems where a perfect decision is not possible, and value judgements will then be needed.

Direct Brain Interface


Present-day trends strongly suggest that direct stimulation of nerves and the interpretation of signals from nerves will be possible by 2010. Artificial (that is, prosthetic) sensors will also exist, opening the possibility of enhanced 'add-on' senses for humans. Furthermore, Reference 1 suggests a full interface between the human brain and a computer by 2025. Whether this is optimistic or even realisable is of course unclear. Certainly it is already possible to control simple actions on a computer simply by thinking (using SQUID (superconducting quantum interference device) technology), although not yet for a computer to stimulate thoughts in the operator. If indeed thought interpretation and stimulation do become possible by 2025, some mind-boggling scenarios arise, some of which have been explored further in Reference 2:

- Human-brain enhancement by linking thoughts in the brain with those in the computer, and using the computer's superior capabilities.
- Linking to a powerful computer may enhance intelligence, consciousness and awareness. Whole new areas of experience would be accessible to the human.
- Telepathy becomes possible, by upload, transmission and download of thoughts.
- By storing a person's mind effective immortality is achieved.
- It would become possible for two people to share the same

Figure 3—Communication via the Medium

consciousness simultaneously, opening up whole new possibilities for personal interaction.

- By uploading physical experiences and sensations, the computer would be able to re-create any life situation for the user. This would give the potential to build alternative reality, the sensation of actually being in an environment with much greater 'reality' than can be offered by mere holograms or conventional virtual reality. This could be linked to other people or simulations.
- Pleasure could be directly created in the brain with no need for any physical effort. This may be addictive!
- Therapy would be revolutionised, as people could be put through actual experiences without danger, in order to gain confidence.
- Punishment for criminals could be revolutionised by the same token; for example, the burglar could experience the trauma of being burgled!
- If the computer is capable of download into the brain, instant training and education becomes possible. Learning and education in the traditional sense would be obsolete since knowledge and skills could be available via the computer. This will tend to 'equalise' people.
- The computer side of human-computer links may develop personalities of their own. It may develop some of the psychological problems suffered by people, such as forgetting, owing to the difficulty in storing and retrieving information in a meaningful way.
- What would the result be if all of these computers are networked while people are linked to them? A global consciousness!



Clearly the above would have profound implications for the individual and humanity. Personalities in the society could be pooled, with artificial personalities being created to add to the options available. Users would be free to choose personality, memories and skills from the common pool, so there would no longer need to be a permanent mapping of mind onto particular brains. The concept of the individual as we know it may cease. For some people, the first time they upload their personality onto the machine and download a another personality may be the last time their own personality is used. They will effectively have died as an individual!

Infrastructure—The Medium

In the Western world of 2025 we might reasonably presume that most individuals will own a computer which by today's standards would be considered immensely powerful. Note however that neither human or machine will exist in isolation. People will interact with each other within society while their computers will similarly interact via the broadband telecommunications environment, delivering high-speed access anywhere on the planet with minimal delay. We will refer to this network of computers and people as the *Medium* (Figure 3).

Within the Medium it becomes irrelevant whether the user wishes to access processing or telecommunications services. Users will simply want to engage in some form of interaction and will not care which resources and mechanisms are invoked on their behalf. With such capabilities it is readily envisaged that the Medium would become a key part of everyday life, business and domestic. It may also develop its own consciousness,

evolving directly from the consciousness that its constituents have; that is, people and their computers. The Medium then becomes an organism created by man.

Work

In the society of 2025, it is unclear what role people would play in running industry and commerce. What is clear, however, is that society is unlikely to keep pace with technological developments. Although menial jobs will generally be done by robots, including skilled manufacturing jobs, some service industries may be more resistant to change. People may still want to be served by real people when they go out for dinner, and have their hair cut by real people. They may even want real nurses in hospital. In the arts, people will be able to explore their own creativity to the full, and there will still be real actors and singers at the theatre. Since I will have retired long before 2025, I may not be too concerned at the lack of demand for real engineers, but others will be!

It is beyond the scope of this article to explore the tensions in society during its evolution. Since most workers may in time be replaced by robots or computers, society would certainly need to restructure, although it would be futile to try and predict its structure, which could be anywhere between an ultra-wealth ultra-poverty mix and a complete socialism. However, assuming that it would lie somewhere in between, with a reasonable income for everyone, we can only speculate how people might relate to the infrastructure.

Personal Devices

Given the power and indispensability of the personal computers

discussed, it seems reasonable to assume that most people would have one. It would be capable of anything we currently dream of having in a computer. The user could communicate with it by natural voice, body language, even by thought if a direct brain link is accomplished. It would play an essential part in the user's life, acting as his agent for just about everything. It would have full access to the Medium. We will call this personal computer a PAL (**P**ersonal **A**gent for **L**ife), it being a descendant of the PDA (personal digital assistant). It can connect to the Medium to offer full functionality, but is still capable of many functions without direct connection. It can read smart cards (by then their capacity should well exceed 100 Tbyte) and the few remaining dumb storage media, so still includes all of the functionality of the electronic book.

Whether the PAL would be the user's only device is uncertain. For example, it may not yet be possible to include all the required functionality into a small, portable device. If not, then the user would still need a larger base at the office or home into which he would plug his PAL to achieve full Medium functionality. The PAL would thus be able to link to the Medium when necessary. Outside, this would rely on mobile communications, making good use of the fixed mobile network too. It would provide access to any communications or processing the user wants.

PALs would store all personal information and preferences, and would be able to act as an agent wherever people need to interface to the outside world, taking care of financial transactions, scheduling, information collection, entertainment organisation etc. People would be liberated to relax and enjoy the finer things in life. Those wanting to escape from it all could go for a walk, leaving their PAL behind to act on their behalf when required. More likely, it would be taken along to act as a guide, interrupting only if

permitted or requested. Being capable of conversation, it would make an ideal companion—talking only when appropriate.

PALs would enhance creativity, by collecting ideas from all over the world and using them to stimulate ideas. This would enable people to develop their full potential. Conversely, since it can carry out most mental work for the user, it may make some people lazy.


The interface would be fully customised for the user, using video display, active lenses, virtual reality or direct brain interface, and responding to any form of input from the user.

Other Technology Leaps

If the Medium truly has the capabilities described, then we could expect that research and development in almost every field would have made huge strides by 2025. In areas limited by theoretical extrapolation and lateral thinking (bringing together superficially unconnected ideas to yield a useful result) such as quantum mechanics and mathematics, we would see the largest gains. Other areas may be limited by the time and resources (or ethics) required to build experiments to find missing data. Indeed, physical resource limits may be a significant problem. Again, some of the following possibilities are explored in Reference 2:

- Many of the crises facing the world at the moment may have been solved in principle, merely awaiting the start of implementation.
- With new solutions for energy, better physics, and better designs, mankind could begin space exploration in earnest.
- New materials would be discovered and new engineering techniques developed. Currently novel technologies such as nanotechnology would mature quickly, possibly leading to picotechnology.
- Safe genetic engineering would have a great impact. Genetic models may allow genes to be designed and simulated before building custom organisms. Many diseases would be eradicated and new medicines invented.
- Life expectancy will increase as a result of better medicine and social care.
- Ethics permitting, parents could specify the properties of their child. Genes could be selected from a gene bank and a custom zygote produced.
- Even if the full brain link is not around by 2025, direct nerve links will be, and devices would be available that provide new 'senses'. We may see tactile devices which sense surface properties much better than our fingers, ears sensitive to a wider range of sound, safety devices which can detect radiation, navigation aides which allow people to 'see' magnetic and electric fields and many others.
- Advanced materials and technology may allow lightweight cybernetic additions to the body to enhance strength or provide protection. Artificial exoskeletons may become more versatile.
- Bodily 'spare part' technology, together with the new range of sensors, could realise the 'bionic man'. Such body 'add-ons', perhaps linked to standard interfaces, may become everyday accessories, allowing people to select and optimise their abilities for the day when they get dressed. (See Figure 4.)
- It may not be too optimistic to expect theoretical understanding of time travel and teleportation, which would confirm whether these are possible in practice and how they could be accomplished if they are.

Figure 4—Potential replacement of body parts as percentage of body weight



Travel

In spite of advances in information technology, and impacts in other technology areas, people will still travel for various reasons. With PALs taking care of the user's travel plans, this activity would become simpler. We will certainly have many smart roads by 2025, and cars will probably be driven automatically on these routes to prevent congestion. This form of travel assumes that teleportation is either not possible, too expensive, or simply hasn't been implemented yet. Hopefully, more efficient use of resources and better energy production would see much of the environment left in good condition (even restored), making travel more enjoyable.

Virtual Reality

Virtual reality will have been around for a long time by 2025. We can thus expect it to be as mature as TV is today. It is likely to be an integral part of the Medium interface, either by direct brain connection or by more conventional interfaces, such as lightweight spectacles, active lenses and large flat screens. Resolution will be indistinguishable from reality and it will be multisensory. It will be much more versatile as a result of Medium capabilities. Even if direct brain interface is not possible, the Medium will inject an extra dimension into interactions. For leisure use, its imagination should be able to produce fantastic alternative realities, with the merest hint of guidance by the operator.

Telepresence

Communication will deliver a 'being there' sensation in line with telesphere principles. This will largely be accomplished using descendants of virtual reality technology, but a direct brain interface would make for an increased integrity. Artificial senses will also be very useful, allowing people to 'sense' an

environment which they could not possibly enter for real, such as the inside of a nuclear reactor. Surgeons would be able to use mini-robots as if they were an extension to their own body.

Leisure telepresence could be very popular, displacing much of environmental impact of real travel, so that people would be able to visit and experience attractive locations without causing erosion and congestion. Social interaction could be enhanced by the same token, especially when we consider the other technologies that will be around. In 2025 cyberspace, there are no language barriers and physics is custom designed!

Social Problems

Since the long-term impact of telecommunications is likely to be greater than that of the industrial revolution, we can expect significant social stresses. 21st century Luddites may resent control of their world transferring to machines, regardless of any benefits, while other people may choose to 'opt out' of the new society, being unable to compete economically. Attempting to control resulting crime could give rise to a big brother society, with electronic surveillance and automatic identification of criminals.

Countries without Medium technology would be at a serious disadvantage, widening the gulf between rich and poor countries. However, rich countries will also be in a better position than ever to help.

There will also be problems for individuals. Short periods of withdrawal into the Medium to escape

from the real world may be commonplace and harmless. However, too much withdrawal may make it increasingly difficult to determine which thoughts are personal and which come from the Medium or other people attached. Absorption and loss of individuality may result. The Medium could offer such potential for enjoyment that addiction and absorption may be common, though some people will be more susceptible than others. People have different mental profiles, including susceptibility, so the Medium interface would need to be customised for each person.

Values

The Medium offers enormous technical potential, but it must be treated with caution as there will inevitably be some problems. It raises many moral, ethical and religious issues, too complex to explore here in detail. The development of consciousness in the Medium could make it susceptible to corruption. We may teach it Asimov's laws of robotics but it may decide to ignore such inconvenient limitations.

Without proper guidance, intelligent, conscious machines might develop value sets quite unlike our own, so it would seem a reasonable precaution to limit the scope of such machines and ensure there is an 'off switch'. The concept of evil machines has often been explored in science fiction, but will soon need to be addressed for real, preferably before it is too late. Even if humans have taught the Medium well, it is bound to make its decisions on a slightly different basis than would humans, which may be a good or bad thing.


Instability

Computer and telecommunications integration can result in potentially unstable complex systems. For instance, in an economic environment, a computer with even a small advantage could quickly become all powerful, but another computer could easily take control with just a small swing in the direction of the source of useful data. Most problems could be avoided by good planning, but there will always be surprises. A brownout of something as critical as the Medium could be catastrophic.

Impact on Telecommunications Networks

Even without advanced computers, network traffic can be expected to grow quickly. But when computers reach the enormous powers predicted, traffic would grow much faster, with the exponent itself increasing. Extra traffic arises from both increased use of the networks by people making use of its vast power, and from the inter-computer traffic that would be the lifeblood of the Medium. However, the Medium would help improve the network infrastructure through its advanced research and development capability. It is very unlikely that the raw network capacity would be exceeded by then. Figure 5 shows schematically how network traffic growth could change through to 2025.

Figure 5—Network traffic growth explosion



We can make an estimate of network digital traffic by 2025. From the earlier discussion we can estimate that the 2025 PAL (personal computer), could have a digital processing capability of 10^{10} MIPS. Assuming a 100 bit instruction width, this gives 10^{18} bit/s internal traffic in the PAL. If even one billionth of this makes it out onto the network, this gives 1 Gbit/s of traffic from each PAL. Fifty million of these in the UK (assuming everyone would have one) would give a total network traffic (from PALs) of 5×10^{16} bit/s. In addition to this we would also have traffic between other computers in the Medium, probably adding about the same as the personal computers. The total network digital traffic could thus be 10^{17} bit/s (one hundred thousand terabit/s), not counting analogue traffic. The Medium would be in a good position to predict network use, because of its intimate relationship with the network users, so would be well placed to do network planning and management.

Conclusions

The world of 2025–2030 will be very unlike that of today. Since 2010, computers would have exceeded Man's intelligence on many fronts and progress would have accelerated. By 2025, the fruits of this astounding progress will have completely transformed computing and telecommunications, merging them into a powerful knowledge and communications infrastructure—the Medium. People would have a very different lifestyle, with a real possibility that work would be a purely voluntary activity. The economy may be completely stable.

As with all futurism, this paper has presented only one possible scenario. However, it is abundantly clear that we must not underestimate the possible effects of continued progress in information technology and indeed biotechnology. Even if reality turns out to be a few orders of magnitude more conservative than

the effects discussed here, there would still be a tremendous impact on mankind.

References

- 1 Systems Research Technology Calendar, 1994.
- 2 ED REGIS. Great Mambo Chicken and the Transhuman Condition. Penguin Books.

Acknowledgements

The author would like to thank Dave Heatley and Peter Cochrane for their valuable inputs and comments.

Biography



Ian Pearson
BT Networks and
Systems

Ian Pearson graduated in 1981 from Queens University, Belfast, with a B.Sc. joint honours degree in Applied Mathematics and Theoretical Physics. After four years at Shorts in Belfast working on missile design, he joined the Performance Engineering Division at BT Laboratories studying data networks and protocols. He has since worked in several divisions, advising on the communication impact of technology developments, especially in computing and information technology. Inventions along the way include a free-format ATM derivative, various novel computer interfaces such as the active contact lens, and many other service and technology concepts. Since 1991, he has worked in the Systems Research Division investigating various aspects of future networks, while keeping a view of developments within IT.

Peter Cochrane and David Heatley

Aspects of Optical Transparency

The optical-fibre amplifier will bring about network transparency and reductions in interface problems, software and operating costs, while improving reliability and performance.

Prologue

The optical-fibre amplifier is set to revolutionise telecommunications on a broad range of operational and application fronts. Simplifications in repeater hardware, the opening up of the latent fibre bandwidth, and the enormous increase in cable capacity are often cited as the major advances heralded by this technology. We should note, however, that these are merely the first rungs of a technology ladder that lead to wider capabilities. For example, wavelength-selective gratings are now being written into the core of optical fibres, with the potential of programmable fibres in the not-too-distant future. These might allow real-time frequency and signal selection within a network in a revolutionary way. Contactless connectors, active leaky feeders and optical wireless are other recent arrivals on the scene that also exploit optical transparency. All-optical communication for mobile units via a fixed transparent optical network is therefore a real prospect.

On a global basis the need for switching remains, but for cities and certain small countries, that requirement could be diminishing as we look towards the early part of the 21st century. Wavelength routing using linear and non-linear fibre properties offers the prospect of non-blocked switchless communication for up to 20 million customers. There is also the prospect of smart fibres with an inherent switching and routing capability.

While the above prospect is stunning when compared with the technology of the previous decade, it is in the area of economics that optical transparency affords a unique oppor-

tunity. What is generally not recognised is that there is a raft of hidden benefits; specifically, major reductions in interface technology, network control and management software; reliability improvements; and reductions in the number of switching sites. System and network operating costs that are a small fraction of that today thus look certain to be realised.

No Hierarchies

The need for thousands of local switches and hundreds of central offices within a nation has necessitated the creation of hierarchical networks where the traffic is thin in the local loop and dense at the core and international level. The various tiers within networks have been organised to assemble and groom the traffic so that it may be most economically transported between customers. The sheer number of switches and transmission systems involved dictates such a topology. However, there is now an alternative. If we exploit the inherent reach capability of optical fibre and extend the local loop boundary to more than 40 km, national networks of less than 100 switches become possible. With such low numbers there is little point in creating a hierarchical structure. In fact it can be argued that a hierarchical structure across the entire planet would be a redundant concept, whereupon the global network becomes purely a local loop: no core, no international, just local!

Economics of Analogue and Digital

History has seen the transmission medium migrate from a single wire

for telegraphy to a twisted pair for telephony, then coaxial for multi-channel telephony and now multiple fibres. In the future, this trend may culminate in a single ubiquitous optically-amplified fibre with varying amounts of active elements in the transmission path. It is likely that with the removal of electronics and other blocks on bandwidth, the resulting photonic transparency will encourage a partial move back to analogue communication! By analogue we refer to those instances which exploit transparency between terminal stations, there being no signal-specific regeneration on route, only broadband amplification. Our definition of analogue therefore spans a blend of analogue and digital formats.

It has to be remembered that in the 1970s the primary reason for going digital was to realise the lowest-cost solution for an international network that was striving to maintain a given quality of communication for telephone traffic. At that time, it turned out that digital switching and transmission gave the lowest-cost solution by about a factor of 2 compared with any combination of analogue and digital (Figure 1). However, this outcome was dictated by repeater spacings, then of about 2 km on twisted pair and coaxial. With repeater spacings now spanning more than 50 km on fibre routes, the economics of networks have changed radically. When this is combined with optical transparency and a more than 100-fold reduction in the number of switching sites, we soon reach a point where a combination of analogue switching and transmission realises the most economic network.

Reliability

With the development of pulse-code modulation in the 1960s, digital transmission was realised, with all the attendant benefits of consistency of signal quality and tolerance to interference. The down side was that large numbers of repeaters were




Figure 1—Network costs with projected technology progression relative to 1970

required between switching centres to reshape, regenerate, and repeat the signals. Early repeaters proved to be the dominant source of system failure, so much so that elaborate measures were introduced to improve end-to-end reliability. Primarily, this involved $N + 1$ stand-by and diversity routing.

$N + 1$ stand-by provides one hot stand-by circuit for every N traffic-carrying circuits within the same cable or duct route. When a failure occurs, the affected traffic is automatically switched over to the stand-by circuit. This offered real benefits at the time of its introduction, but detailed studies have shown that this is no longer the case, particularly since the introduction of optical fibre. Figure 2 shows that by duplicating the power supplies associated with repeaters and other active elements (bearing in mind how few these elements are becoming), circuit availability is improved to a level on a par with $N + 1$ stand-by, particularly

on long lines, but at a fraction of the cost in terms of equipment, maintenance and network overheads.

Diverse routing differs only in one key aspect from the $N + 1$ strategy: the stand-by route follows a different physical path between the common end points, which enables a better link availability to be achieved. The geographic separation between the main and stand-by routes can range from a few kilometres within cities to hundreds of kilometres within countries, and ultimately continents for international links. Again studies have shown that the original rationale behind diverse routing is no longer valid, with power supply duplication producing a similar reliability improvement. However, since diverse routing is generally limited to long-distance high-capacity links because of its high implementation and operating cost, and since it affords a useful level of network flexibility, there is no compelling argument to see its wholesale




Figure 2—Unavailability of 100 km, 1000 km and 10 000 km systems with or without $N + 1$ stand-by

removal. In any case, there is a general trend towards self-organisation within networks that reduces the need for diverse routing while at the same time improving plant utilisation.

Software

Today's networks rely heavily on software for their management and service provision. In contrast to




Figure 3—Software and hardware reliability trends

hardware, where failures generally have a small localised impact, minor errors in software pose a considerable and widespread risk. If the present trajectory in software development is maintained, the magnitude of the risk will grow exponentially as we look to the future. Indeed, we are already witness to the reliability of hardware improving rapidly while that of software is reducing (Figure 3), leading inevitably to sub-optimal system and network solutions. From any engineering perspective, this growing imbalance needs to be addressed. The introduction of optical transparency is likely to see a reduction in the scale and complexity of software through the corresponding reductions in switching and routing.

Critical Mass

Because modern networks contain thousands of nodes, the effect of individual node failures tends to be localised and isolated—barring software-related events! Furthermore, the impact of single or multiple failures is effectively governed by the 'law of large numbers', with individual customers experiencing a reasonably uniform grade of service. However, as the number of nodes is reduced, the potential for catastrophic failure increases, with the grade of service experienced at the periphery becoming extremely variable. The point at which such effects become apparent depends on the precise network type, configuration, control and operation, but as a general rule networks with less than 50 nodes require careful design to avoid quantum effects occurring under certain traffic conditions. That is, a failure of a node or link today for

a network based solely on fibre with less than 100 nodes may be expected to have a fault rate of less than 20 per cent that of today's networks

a given network configuration and traffic pattern may affect only a few customers and pass almost unnoticed. The same failure tomorrow could be catastrophic, affecting large numbers of customers, merely because a different configuration and traffic pattern existed at the time.

Network Management

Monitoring the operation of systems and networks, extracting meaningful information and taking appropriate action to maintain a given grade of service are becoming increasingly complex and expensive. The level of complexity increases in proportion (at least) to the amount of data being handled. Consider for example the quantity of data generated by the failure of a switching node. That node generates a fault report while each of its related nodes generates an error report. For a fully interconnected network of N switching nodes, this means one failure report plus $N - 1$ error reports. If in a large network we allow for two or more nodes failing simultaneously, then, for example, a network of 500 000 nodes, each with a mean time between failures (MTBF) of 10 years will suffer an average of 137 node failures and will therefore generate an average of 68.5 million reports per day. This assumes that each node is communicating with all the others, which is of course somewhat extreme. However, even if we consider the least connected case, we find that the mean number of reports per day

$$\approx \frac{(N^2/6)}{\text{MTBF in days}}$$

which shows that the reports still count in the millions. Indeed, a large national network with thousands of switching nodes can generate information at rates of about 2 Gbyte/day under normal operating conditions. Clearly, maximising the MTBF and minimising N must be key design objectives in reducing the information mountain.

Human Interdiction

It is important to recognise that network reliability is as much a people issue as it is to do with equipment and technology. For example, it is estimated that over 50 per cent of the faults experienced by telcos today are in the local loop, around half of which are in some way attributable to their own peoples' activities. Replacing manual distribution points and wiring frames with software-based routing over a passive optical network (PON) thus represents a significant gain. When combined with gains arising from the eradication of corrosion, the overall saving in the local loop can be more than 40 per cent of the operating total. Similar savings are also possible for repeater stations and switches. With all things considered, a network based solely on fibre with less than 100 nodes may be expected to have a fault rate of less than 20 per cent that of today's networks. This figure assumes fibre to the home (FTTH). If however the fibre to the kerb (FTTK) option is adopted, with copper or radio drops to the customer, the above gains will not be realised and the overall running costs will remain high.

An Industry Tension

It is interesting to reflect that telcos think in terms of switches while the computer industry thinks in terms of routers. On the one hand, the computer industry has a desire for total freedom and anarchy in communications with no one intervening in the process. On the other hand, the telcos wish to have switches at strategic locations so they can manage the overall process. The two halves of the IT industry are thus diametrically opposed. In reality, a mix of both philosophies will be required in networks for some time. To opt for either as a total solution would merely lead to a sub-optimal solution for a sizeable user population.

Technology Speed-Up

All of this is compounded by the fact that technological development is accelerating to the point where most standards are dead on arrival, or at least moribund, and have little hope of being universal. Yet we have to find a solution that allows us to inter-work on a global basis. Optical transparency has a key role to play in this context. If customers can be allocated a unique (or shared) carrier and be left to utilise it as they see fit, then for the first time the telecommunication industry can stand back and watch the interface war rage. Bluntly, users don't care whether they are communicating via MPEG, JPEG, X.25, X.400 or any other standard: all they care about is initiating the call/connection and having their terminal equipment and the network do the rest. Significantly, some computer manufactures are already installing software programmable codecs in their products to enable them to communicate with all forms of modem, voice and vision codecs as they emerge.

Closing Remarks

The Clinton/Gore initiative for a national super-highway is being mirrored in Europe under the 4th framework programme and is, to some extent, in place in the UK in the form of the SuperJANET network. Opening up high-speed telecommunications in this way will cause the computer industry to release machines of a very different nature than hitherto. Perhaps we will at last see an end to excessive code compression on visual and speech signals. Perhaps we will be able to realise distributed computing and employ telepresence, teleconferencing, multimedia and virtual reality as a means of communication within a humanised framework that removes the need to travel. Without transparent optical networks none of this will be technologically nor economically achievable.

Bibliography

COCHRANE, P., HEATLEY, D. J. T., *et al.* Optical communications—future prospects. *IEE Electronics & Communication Engineering Journal*, Aug. 1993, 5(4), pp. 221–232.

COCHRANE, P. and HEATLEY, D. J. T. Reliability aspects of optical fibre systems and networks. *IEEE International Reliability Physics Symposium*, San Jose, USA, 11–14 April 1994.

HILL, A. M. Network implications of optical amplifiers. *Optical Fibre Communication Conference (OFC'92)*, San Jose, USA, 2–7 Feb. 1992, paper WF5, p. 118.

Modelling Change in Telecommunications. *BT Technol. J.*, Special Issue, April 1994, 12(2).

Biographies



Peter Cochrane
BT Networks and
Systems

Peter Cochrane graduated from Trent Polytechnic with a B.Sc. Honours degree in Electrical Engineering in 1973 and gained an M.Sc., Ph.D. and D.Sc. in Telecommunications Systems from the University of Essex in 1976, 1979 and 1993 respectively. He is a fellow of both the IEE, IEEE and Royal Academy, a visiting professor to Essex and Southampton Universities and an honorary professor at Kent. He joined BT Laboratories in 1973 and has worked on both analogue and digital systems. He has been a consultant to numerous international companies on projects concerned with systems, networks and test equipment. In 1978 he became manager of the Long Lines Division which received the Queen's Award for Technology in 1990 for the production of optical receivers for TAT-8 and PTAT-1 undersea cable systems. In 1991, he was appointed to head the Systems Research Division at BT Laboratories which is concerned with future advanced media, computing and communications developments. During 1993 he was promoted to head the Research Department at BT Laboratories with 620 staff dedicated to the study of future technologies, systems, networks and services. He is also the Development and Procurement board member for technology.



David Heatley
BT Networks and
Systems

David Heatley obtained his doctorate in Optical Communications Systems from the University of Essex in 1989. He joined BT Laboratories in 1978 to work on the development of analogue and digital optical fibre systems designed specifically for video services. In 1985 he was appointed head of a group responsible for the development of optical receivers for terrestrial and undersea applications. In this capacity he was a member of the team that received the Queen's Award for Technology in 1990. He is now with the Advanced Media Unit and heads a team with special responsibility for future technologies and telemedicine. During his career he has published widely on telecommunications, ranging from discrete components, through systems to networks. He is a Member of the IEE and a Chartered Engineer.

Peter M. Adams and Nigel Wall

Global System for Mobile Communications

The Development of the GSM Standard

The introduction of analogue cellular radio telephone 10 years ago resulted in a rate of expansion of users that far exceeded market expectations. During this period, the European digital Global System for Mobile Communications (GSM) standard has been developed. Initially introduced only 2 years ago, the take-up has been exceptional. The GSM standard, that many believed would fail because of its complexity, is already proving to be one of Europe's most successful standards.

Introduction

Mobile telephony systems have been available for about 40 years. Early systems were expensive vehicle-mounted sets that addressed the needs of a small market. The introduction of analogue cellular technology in the early-1980s stimulated the present expansion of service to lightweight handsets. Ten years later, we are witnessing the expansion of the next stage, which will enable mobile telephony to reach a mass market, predicted by some to approach that of today's public switched telephony network (PSTN). The development in Europe of the digital Global System for Mobile Communications (GSM) system, and its offshoot Digital Communication System (DCS 1800) started in the 1980s, and entered operation in various European countries from 1992 onwards. This article reviews the background to the development of the standard.

The Development of Cellular Systems

The first mobile telephone networks were manually operated and the terminals were heavy, bulky and expensive. The service area was restricted to the coverage of a single transmission and reception site (single-cell systems). Little radio spectrum was available for this kind of service, because of the demand from military systems and broadcasting (in particular television). As a consequence, the capacity of the early systems was small, and saturation

came quickly despite the high cost of terminals which deterred many potential customers. Quality of service decreased rapidly with congestion, and the throughput sometimes fell drastically due to near-deadlock situations.

During the 1970s, large-scale integration of electronic devices and the development of microprocessors opened the door to the implementation of more complex systems. Because the coverage area of one antenna is limited mainly by the transmitting power of mobile stations, systems were devised with several receiving stations for a single transmitting station. They allowed coverage of a large area at the cost of additional infrastructure complexity. But the real breakthrough came with cellular systems, where both transmitting and receiving sites are numerous and with individual cell coverage areas that partially overlap. Interference is avoided because adjacent sites operate on different frequencies.

Cellular systems also employ the concept of frequency re-use: the same frequency is used by several sites, or cells, each time for a different user, resulting in a tremendous gain in system capacity. The counterpart is the increased complexity, both for the network and for mobile stations, which must be able to select a frequency channel among several possibilities.

These analogue systems are referred to as *first-generation mobile systems*. Several incompatible systems were developed and used in Europe. In the UK, TACS (Total

Access Communication System) was developed from AMPS (American Mobile Phone System) and entered service in 1985. Most European countries have one or more cellular networks today. Table 1 shows the major cellular networks in operation in Europe at the end of 1994. The extraordinary rate of adoption of GSM over two years is evidence of the success of the standard.

The first-generation cellular systems are based on analogue speech transmission with frequency modulation. They use frequency bands around 450 MHz or 900 MHz. Their coverage is usually nationwide and their capacity reaches several hundreds of thousands of subscribers. The largest national systems in Europe are the two British TACS operations provided by Cellnet and Vodafone, each with about one and a half million subscribers at the end of 1994. The highest population penetration is held in Scandinavia, with more than 13% of the Swedish and Norwegian population having mobile equipment. These figures are much higher than the mean European values. For instance, the penetration factor in 1994 in France was only about 1.5%, with the UK at around 6%.

Mobile equipment evolved very rapidly during the late-1980s. At the onset, only vehicle-mounted equipment could be built. In the mid-1980s, portable equipment appeared, weighing a few kilograms and with a recharge battery life of a few hours. Handheld equipment first appeared around 1988, not small enough to fit in a pocket, but fitting nicely in an attaché-case. In 1994, the smallest terminals on the market weighed less than 200 g, and fitted in a shirt or trouser pocket.

Why Have a New Standard?

One key objective was that the new system must allow free roaming of the subscribers within Europe. Thus a subscriber of a given national network may access the service provided by another operator when

Table 1 Major Cellular Systems in Europe at the End of 1994

Country	Systems	Frequency Band	Launch Date	Subscribers (thousands)
United Kingdom	TACS	900	1985	3100
	GSM	900	1992/4	140
	DCS	1800	1993/4	330
Scandinavia (Sweden, Norway, Finland, Denmark)	NMT	450	1981	670
	NMT	900	1986	1560
	GSM	900	1992/3	910
France	Radiocom 2000	450, 900	1985	280
	NMT	450	1989	144
	GSM	900	1992	460
Italy	RTMS	450	1985	16
	TACS	900	1990	2150
	GSM	900	1992	75
Germany	C-450	450	1985	720
	GSM	900	1992	1700
	DCS	1800	1994	30
Switzerland	NMT	900	1987	290
	GSM	900	1993	37
The Netherlands	NMT	450	1985	21
	NMT	900	1989	230
	GSM	900	1994	68
Austria	NMT	450	1984	46
	TACS	900	1900	220
	GSM	900	1993	13
Spain	NMT	450	1982	34
	TACS	900	1990	380

travelling abroad. The same GSM mobile station must enable its user to call or be called anywhere within the international coverage area.

As early as 1982, the basic requirements for GSM were stated. They were slightly revised in 1985, and are still largely current—reproduced verbatim in Panel 1 (overleaf). The most significant change is that the focus is on handheld mobile stations, rather than transportable or vehicle-mounted stations.

These system objectives were agreed and distributed to the telecommunications industry, with the objective of providing a clear framework for progressing the technical design. It was also clear that spectrum efficiency and hence the capacity offered by the system should be better than the existing analogue networks.

The Creation of GSM

A major prerequisite for a common radio system is a common radio

spectrum allocation. Negotiation of spectrum allocation is a complex and lengthy business. However, the needs of cellular systems were identified in good time and in 1978 it was decided to reserve a frequency band of twice 25 MHz at around 900 MHz for mobile communication in Europe.

A second requirement is that all equipment must be built to comply with a common standard. The CEPT (Conférence Européenne des Postes et Télécommunications) is a standardisation arena which, in the early-1980s, included the European administrations of posts and telecommunications of more than 20 countries. CEPT established a new standardisation project group to specify a unique radio communication system for Europe, at 900 MHz. The new-born *Groupe Spécial Mobile* (GSM) held its first meeting in December 1982 in Stockholm, under the chairmanship of Thomas Haug, from the Swedish Administration.

Panel 1—Basic Requirements Set Out By GSM (original text as written by the GSM committee in 1985)

Services

- The system shall be designed such that mobile stations can be used in all participating countries.
- In addition to telephone traffic, the system must allow maximum flexibility for other types of services, e.g. ISDN related services.
- The services and facilities offered in the PSTN/ISDN and other public networks should as far as possible be available in the mobile system. The system shall also offer additional facilities, taking into account the special nature of mobile communications.
- It should be possible for mobile stations belonging to the system to be used on board ships, as an extension to the land mobile service. Aeronautical use of GSM mobile stations should be prohibited.
- In addition to vehicle-mounted stations, the system shall be capable of providing for handheld stations and other categories of mobile stations.

Quality of service and security

- From the subscriber's point of view, the quality for voice telephony in the GSM system shall be at least as good as that achieved by the first generation 900 MHz analogue systems over the range of practical operating conditions.
- The system shall be capable of offering encryption of user information but any such facility should not have a significant influence on the costs of those parts of the system used by mobile subscribers who do not require such facility.

Radio frequency utilisation

- The system concept to be chosen shall permit a high level of spectrum efficiency and state-of-the-art subscriber facilities at a reasonable cost, taking into account both urban and rural areas and also development of new services.
- The system shall allow for operation in the entire frequency band 890–915 MHz and 935–960 MHz.
- The 900 MHz CEPT mobile communications system must co-exist with earlier systems in the same frequency band.

Network aspects

- The identification plan shall be based on the relevant CCITT Recommendation.
- The numbering plan shall be based on the relevant CCITT Recommendation.
- The system design must permit different charging structures and rates to be used in different networks.
- For the interconnection of the mobile switching centres and location registers, an internationally standardised signalling system shall be used.
- No significant modification of the fixed public networks must be required.
- The GSM system shall enable implementation of common coverage PLMNs.
- Protection of signalling information and network control information must be provided for in the system.

Cost aspects

- The system parameters shall be chosen with a view to limit the cost of the complete system, in particular the mobile units.

Thirty-one people from 11 countries were present at this first meeting.

In 1990, at the request of the UK Department of Trade and Industry, the specification of a version of GSM adapted to the 1800 MHz frequency band was added to the scope of the standardisation group, with a frequency allocation of twice 75 MHz. This variant, referred to as *DCS 1800* (Digital Cellular System 1800 MHz) is aimed at reaching higher capacities in urban areas for the type of mass-market approach known as *PCN* (personal communications network). BT was actively involved in this part of the development until the application for a PCN license was turned down.

The elaboration of the GSM standard took almost a decade. Major milestones are shown in Table 2, and the corresponding stages are described in more detail in the

Table 2 GSM Project Milestones

Date	Achievement
1982	<i>Groupe Spécial Mobile</i> is created within CEPT
1986	A Permanent Nucleus is set up
1987	Main radio transmission techniques are chosen, based on prototype evaluation (1986)
1989	GSM becomes an ETSI technical committee
1990	The Phase 1 GSM 900 specifications (drafted 1987–1990) are frozen. DCS 1800 adaptation starts
1991	First systems are running (Telecom 91 exhibition,...). DCS1800 specifications are frozen
1992	Many major European GSM 900 operators begin commercial operations
1993	Roaming agreements between several operators put into practice
1994	Data capabilities launched

Note: The span of GSM work from the very start to commercial service has extended over some 12 years, but the actual specification work did not start until 1987, and work continues today

Figure 1—Structure of committees in ETSI involved with GSM

following text. GSM is considered to be a second-generation mobile system. Although adopted in many parts of the world, it is not the only second-generation mobile system. Different standards have also been developed within North America and Japan.


Organisation of the Work

The first two years of the GSM were dedicated to discussions of fundamental principles. The frequency of meetings and the number of participants increased steadily. At the beginning of 1984, three working parties were created; the meetings were split to allow more in-depth technical work. The extent of contributions and problems to solve increased steadily, and at the end of 1985, it became obvious that the number of meetings was insufficient. It was then decided (for the first time in CEPT) that the working parties would meet independently, reporting to the GSM plenary meetings which still endorsed the decisions. The role of the different working parties had already been clearly established for some time:

- WP1 (Working Party 1), for the definition of services;
- WP2 (Working Party 2), for the specification of radio transmission; this subject stayed dominant until 1987; and
- WP3 (Working Party 3), for all other issues; that is, mainly network architecture, specification of the signalling protocols and of the open interfaces between network entities.

Later, a sub-group of WP3 was spun-off as a fourth working part (WP4) to deal specifically with the implementation of data services.

Other working parties were commissioned as sub-technical committees as the system definition became more specialised.



During 1985, a detailed list of recommendations to be produced by the group was discussed and agreed. This followed the model of the Technical Recommendations of the CCITT (Comité Consultatif International Télégraphique et Téléphonique). From 1986 onwards, work became centred on drafting these recommendations. The list includes more than 130 recommendations for Phase 1, sorted in 12 series with a total of more than 30 000 pages! A detailed action plan was generated to follow the progress of this huge task, and was updated at each meeting. These recommendations cover the full specification of the radio interface (that is, the interface between the mobile stations and the infrastructure) as well as fairly detailed specifications of infrastructure architecture and of some of the interfaces and signalling protocols between network entities. For a description of the architecture see Reference 2.

In 1988, the European Telecommunications Standards Institute (ETSI) was created and most of the CEPT technical standardisation activities were transferred to this new body, including GSM. Unlike CEPT, ETSI is not restricted to operators and administrations, but also includes members from industry and user groups. GSM had already anticipated this need by officially allowing industry representatives to participate directly to the working parties on an ad hoc basis since 1987.

With the transfer to ETSI, most GSM Recommendations became I-ETSS (Interim European Telecommunications Standards), then ETSSs. Becoming an I-ETS or ETS requires several stages of approval, including public enquiries and voting, and this process takes several months. In the meantime, GSM Recommendations are called *GSM Technical Specifications* (GSM TS).

At the end of 1991, the GSM Technical Committee was renamed SMG (Special Mobile Group), with WPs renamed as Technical Sub-Committees (STCs), with the same suffix number. The scope of work was increased to include post-GSM developments with Technical Sub-Committee SMG 5 dedicated to third-generation mobile standards, which in ETSI are also known as *UMTS* (Universal Mobile Telecommunication System) and in ITU-T as *FPLMTS* (Future Public Land Mobile Telecommunication Systems). The current structure is shown in Figure 1, which shows the involvement of other expert groups within the ETSI structure. The method of specifying services and facilities developed by NA 1 was adopted, and the definition of the enhancement to Signalling System No. 7 is shared with SPS 2. The original permanent nucleus project team PT 12 has been assisted by several other project teams looking at specific aspects of the onward development of the GSM standard.

The GSM MoU

In parallel with the drafting of technical specifications within the GSM committee, European public telecommunication operators (most of them GSM-operators-to-be) recognised the importance of cooperation for commercial and operational aspects, and signed a Memorandum of Understanding (MoU) in Copenhagen, on 7 September 1987. This memorandum, generally referred to as the *GSM MoU*, covers areas such as time-scales for the procurement and the deployment of the system, compatibility of numbering and routing plans, 'concerted' service introduction, harmonisation of tariff principles and definition of accounting procedures. The memorandum was later signed by more operators, and amended in mid-1991 to accept members from non-CEPT countries and to extend its scope to cover cooperation agreements with non-signatory bodies. By early 1992, in most western Europe, one or more network operators had signed the GSM MoU. In addition to operators, regulatory bodies also signed the GSM MoU, such as the DTI (Department of Trade and Industry) in the UK. DCS 1800 operators have their own association, which has a close working relationship with the GSM MoU.

From 1990 onwards, GSM started to spread outside Europe: a number of countries, such as The United Arab Emirates, Hong Kong, New Zealand and Australia, envisaged adopting the GSM standards. In 1992, Australian operators officially became the first non-European signatories of the GSM MoU.

The Approach to Developing the Standard

General trends in the telecommunication arena made it likely that the system would be based on digital transmission, and that speech would be represented by a digital stream at a rate of the order of 16 kbit/s. The

official decision, however, was not made until 1987.

From 1984–1986, GSM focused on the means to compare different technical possibilities for transmission (digital or analogue), in particular with regards to their respective spectrum efficiency. It was decided to compare several technical proposals on the basis of working prototypes. The results of the comparison were reported at the beginning of 1987. The discussions were difficult, because of the prestige and the advance that would be conferred to a proponent should one of the proposals be adopted. To circumvent this problem, none of the proposed solutions was selected! Only the key features of the transmission method were decided: they were summarised in what was called the *broad avenue*. The narrowing would have to take place afterwards, within GSM WP2. The key features were the following:

- medium-mixed band (200 kHz carrier separation), to be compared to narrowband systems (12.5 kHz or 25 kHz as in existing analogue systems) or wideband systems (one of the candidates proposed a 6 MHz carrier spacing);
- digital speech transmission at a rate not exceeding 16 kbit/s;
- time multiplexing of order 8, with future evolution towards multiplexing of order 16 once a second-generation speech codec has been defined at a smaller rate;
- slow frequency hopping capability.

The GSM Technical Specifications (TSs)

It was thought that speech-encoding techniques would evolve quickly. The plan was to have a first speech encoding algorithm at around 16 kbit/s (13 kbit/s in practice), and a few years later, a more efficient scheme using half as many bits (the latter would be known as the *half-*

rate speech codec). About twice as many subscribers could then be accommodated within the same spectrum allocation, with a radio channel for carrying 13 kbit/s speech (a full-rate channel) being replaced by two radio channels to carry speech encoded with the future algorithm. The whole transmission system was designed from the start to support migration to the future half-rate speech coding although a definite date for its introduction was not determined.

The initial idea, of completing all specification work before the validation phase, was overtaken by the concept of a phased development of an evolving standard. The idea of a Phase 2 step as a functional enhancement of Phase 1 gradually crystallised. Around 1988, the idea matured, and it was agreed that the launch in 1991 would not be with the full range of services. Services were split between two phases; Phase 1 was limited to the most common ones (such as call forwarding and call barring), and alternate services (for instance, speech alternating with data within the same communication). Several supplementary services and the capability of offering data services on half-rate channels were transferred to Phase 2, though all the required specifications existed.

The GSM standard will continue to evolve because of increasing customer expectations and the ever accelerating rate of technical evolution. Future very large-scale integration (VLSI) and microprocessor performances will allow increasingly more complex systems. By the middle of 1991, the basic specifications to enable GSM Phase 1 networks and equipment to go ahead were in place. The emphasis on the work in ETSI on GSM then shifted towards producing the Phase 2 specifications.

GSM Phase 2

Originally it was hoped to complete the Phase 2 standard within 12 months of the completion of Phase 1

in 1991. This would have meant that the standards would have been available about the time public GSM networks were ready to be launched. However, it soon became clear that the amount of work involved in specifying supplementary services, together with the modifications and extensions required to the mobile network specifications for Phase 2, had been seriously under-estimated.

As companies started to implement the Phase 1 standards, errors and improvements to the original specifications were found. Hence a

large number of change requests needed to be processed before the standards could be considered technically stable. This has been a common experience with many other major international standards and recommendations; X.25, for example, was first published by the CCITT in 1976, but major modifications were still being made to the recommendation in 1984—eight years later! It is not surprising therefore that it has taken until 1994 to get the GSM standards anywhere near frozen, and large numbers of change requests on certain specifications (for example, on GSM TS 09.02—the Signalling System No. 7 Mobile Application Part Version 2), were submitted to the SMG plenary as late as October 1994. Amendment requests are still being received for the GSM standards.

A major issue for GSM is cross-phase compatibility. What happens if a GSM mobile terminal built to one phase of the specification roams onto a mobile network built to a different phase? Difficult decisions had to be taken to assess whether a particular feature of Phase 1 had to be retained, or whether a degree of incompatibility might be allowed in order to improve significantly the functionality of Phase 2. A number of incompatibilities between the GSM Phase 1 and Phase 2 specifications were allowed. To minimise problems in this area, the following philosophy was adopted:

- GSM Phase 1 should be known as the *Introductory Phase*.
- GSM Phase 2 should be regarded as the *Main Phase*, with the intention of ultimately updating any Phase 1 systems to Phase 2.
- The next phase would be known as *GSM Phase 2+* (not GSM Phase 3) to emphasise the point that from now on all additions/modifications must be strictly backwards compatible.

A list of the extra services/supplementary services for GSM

Phase 2 is contained in Table 3 (a later article will explain the nature of the teleservices and supplementary services used). The objective has been to achieve convergence with ISDN service definitions, but in a few cases, such as on the closed user group, there are some differences. In these cases, the GSM community felt strongly that a different service definition was more appropriate for a mobile network.

Another important feature of GSM in providing mobility management is the provision for smart cards on GSM handsets. In GSM this is known as the *subscribers identification module* (SIM). Through the use of the SIM, other opportunities are opened up such as roaming between DCS 1800 and GSM 900 networks, and interactions between GSM customers and fixed network installations (assuming a common smart card design for fixed and mobile networks can be agreed).

The Half-Rate Speech Codec

As stated earlier, an important objective for GSM Phase 2 was to develop a half-rate speech codec to double the available capacity in a given band for GSM users. This has proved the most difficult problem to solve for Phase 2. The problems have included the following:

- The speed and the complexity of the microprocessing required needed a greater advance in available technology than was originally realised.
- Even the GSM full-rate codec has not yielded speech quality as good as originally hoped, and any further significant degradation was considered unacceptable.
- There were lengthy disputes on the rules for the process for selecting first the short list of codec candidates, and then the winner of the contest.

Table 3 GSM Supplementary Services up to Phase 2

1	Calling line identification presentation (CLIP)
2	Calling line identification restriction (CLIR)
3	Connected line identification presentation (COLP)
4	Connected line identification restriction (COLR)
5	Malicious call identification
6	Call forwarding unconditional
7	Call forwarding on mobile subscriber busy
8	Call forwarding on no reply
9	Call forwarding on mobile subscriber not reachable
10	Call transfer
11	Mobile access hunting
12	Call waiting
13	Call hold
14	Completion of calls to busy subscribers
15	Three party service
16	Conference calling
17	Closed user group
18	Advice of charge
19	Freephone service
20	Reverse charging
21	User-to-user signalling
22	Barring of outgoing calls
23	Barring of outgoing international calls
24	 incoming calls directed to non-CEPT countries
25	 excluding those directed to the HPLMN country
26	 outgoing calls when roaming outside the HPLMN
27	Barring of incoming calls
28	 incoming calls when roaming outside the HPLMN

Table 4 Phase 2+ Work Item Status List

Title	Completion Date†	Title	Completion Date†
Advice of charge on special services (AoCSS)	Q1/95	Provisions for hot billing	
BSS repeaters (RPT)	Completed	Radio local loop (RLL) using GSM	Q4/94*
Call deflection	Q1/95	Service to GSM handportables in trains	Q1/95
Call forward enhancement (CFE)	Q4/96	SIM data download	
Completion of calls to busy subscriber (CCBS)	Q4/94	Software version number	
Completion of calls to subscriber when no reply (CCNRy)	Q4/94	Standards access SMS, SC-SMES	Q2/95
Completion of calls when subscriber not reachable (CCNRc)	Q4/94	Subscriber class for AoC	
Compression of user data	Q2/95	Support of centralised data interworking function	Q2/95
DCS 1800 4 watt mobile power class	Q4/94	Support of home area priority	
DECT access to GSM networks	Q2/94	Support of optimal routing (SOR)	Q2/95
Direct subscriber access and restriction (DSAR)	Q1/95	Support of operator specific services when roaming (CAMEL)	Q3/95
Enhanced full rate speech codec	Q4/96	Support of private numbering plan (SPNP)	Q3/95
Explicit call transfer (ECT)	Q3/95	Technical enhancements/improvements for phase 2	Q3/95
Extensions to the SMS alphabet	Q1/95	Three volt technology SIM	Q4/94
Facsimile enhancements	Q1/95	Universal access to freephone numbers	Q1/95
Fast moving mobile station	Q3/94*	Update of ME list of network names over the air interface	
General packet radio service (GPRS)	Q1/96	UPT Phase 1	Q3/95
GSM-DEC roaming	Completed	UPT Phase 2	Q4/95
Handsfree MS		User-to-user signalling	Q1/95
High speed circuit switched data	Q1/96	Voice broadcast service (VBS)	Q2/94*
IMEI check digits		Voice group call service (VGCS)	Q2/94*
Interworking with non-GSM applications on the SIM	Q2/95	<i>The following work items were agreed at SMG Meeting No. 11 (Dusseldorf, June 1994)</i>	
Interworking with satellite mobile systems	Q2/96	Enhanced full rate speech codec	
Location services	Q2/95	Handsfree MS	
Malicious call ID (MCID)	Q2/95	High speed circuit switched data	
Mobile access hunting (MAH)	Q2/95	IMEI check digits	
Multiple subscriber profile (MSP)	Q3/95	Overload control enhancement for location registers	
New barring services (NBS)	Q4/95	SIM data download	
Operation of dual band GSM/DCS by a single operator	Q3/95	Software version number	
Overload control enhancement for location registers		Special SMS message indication	
Payphone services	Q1/96	Update of ME list of network names over the air interface	
Premium rate services	Q1/95		
Priority set-up service (PSUS)	Q2/94*		

† Best estimate at time of writing

* The completion date for this work item applies to the preliminary work required to produce the report or requirements. This completion date may be extended (if necessary) to cover subsequent work to produce any technical standards (stages 1, 2, and 3 and delta specifications) which could result. The completion date given to any other work items applies to the completion of the whole work item, including all documentation (stages 1, 2 and 3 and any delta specifications).

● In this area it is naïve to think that technical choices could be made without considering the implications on intellectual property rights (IPR). This has caused problems of procedure in ETSI because agreement could not be reached to implement an IPR policy. Hence there was no agreed ETSI way of dealing with the impact of IPR issues on technical systems to be chosen.

Version 3.3 of the codec from Motorola was finally selected in Autumn 1994.

Data Services and GSM Phase 2+

During the development of Phase 2, an important feature of the GSM

standards has been the growth in data capabilities. In about 1991, the initial version of the GSM short message service (SMS) was completed. Since then, the SMS service has been further developed and has many potential applications. One example is notification of voice, fax or electronic mail information held by a network-based call completion system. Paging replacement would also be possible.

Two versions of a facsimile service have been specified: transparent and non-transparent (that is, error corrected). Transparent operation allows any transmission errors to appear on the facsimile output, whereas non-transparent operation causes packets with errors to be retransmitted and an error-free copy

to be produced. The problem with the latter service is that the facsimile protocol allows only a very short delay due to error correction; the introduction of a significant number of errors would cause the system to time-out and clear the call before transmission is complete. The two versions of the fax service perform differently under different types of network error conditions, and the market has yet to decide which is the overall favourite.

ETSI SMG is currently working on standards for GSM Phase 2+. A set of the new services and features being specified is shown in Table 4. Data capabilities are being further extended to include standards on the generalised packet radio services (GPRS), offering instantaneous data

rates of up to 9.6 kbit/s. These are expected to be available in the 1996/97 period. There are also proposals to concatenate time-slots to provide capacity in multiples of 9.6 kbit/s. Other features being developed include dual-band mobiles, which would enable users on PCN (DCS 1800) networks in the UK and Germany to be able to roam onto GSM networks in Europe. There is also a drive to increase compatibility between the GSM standards and the intelligent networks (INs) being developed on the fixed network.

The Development of GSM Networks

Most major European countries now have two or more live GSM/PCN operators. Increasingly, GSM is also being used outside Europe in places such as Russia, the Middle East, Australia and the Asia-Pacific region. The growth of mobile traffic everywhere is rapid. However, it has yet to be proven whether the current speech quality on GSM is good enough to enable mobile to become a major carrier of traffic substituted for fixed network calls. The audio performance of mobile-to-mobile calls is particularly vulnerable. At present most of the cellular traffic is ancillary/additional to fixed network traffic, in applications where the user is on the move, away from easy access to fixed telephones.

GSM was designed from the outset to be a pan-European system, and in recent years it has become clear that it will become near-global in coverage, including some parts of North America. An important feature of the development of GSM networks has been the development of bilateral roaming agreements between different network operators. The extent of Cellnet's present roaming arrangement are described in Reference 3.

The Future

Further extensions of the GSM standards may include an interworking capability with satellite

systems, a redesigned radio access protocol to include a broadband data capability, and specifications to enable interoperability with the US personal communications services (PCS). Thus, even now, the process of writing GSM standards is far from complete!

The current debate in ETSI is going back to fundamentals. If all the suggestions in the previous paragraph were implemented, would the system still be GSM? Some refer to these latest proposed developments as *Generation 2.5*. Alternatively, should GSM come to an end with the Phase 2+ development programme to allow a period of stability for pay-back on investment? At a later point in time, a jump could then be made to third-generation mobile systems. These decisions on strategic direction have yet to be taken. What is clear, however, is that the GSM systems are being implemented both inside and outside Europe, with exceptional rates of take-up, and that the standard must be considered to be a major success story.

References

- 1 MOULY, MICHEL, and PAUTET, MARIE-BERNADETTE. The GSM System for Mobile Communications. Available from Michel Mouly, enquiries to Fax: +33 1 69 31 03 38.
- 2 BRYDON, ALASTAIR. Global System for Mobile Communications. An Introduction to the Architecture. *Br. Telecommun. Eng.*, Jan. 1995, **13**, p. 287.
- 3 HALL, KEN. Global System for Mobile Communications. A foundation for Successful Competition, Consumerisation and Convergence. *ibid.*, Jan. 1995, **13**, p. 281.

Acknowledgement

The authors of this article wish to acknowledge the cooperation and help of the authors of the book *The GSM System for Mobile Communications*—

Michel Mouly and Marie-Bernadette Pautet. In particular, thanks are given for permission to base part of this article on extracts from the opening chapter of the book¹.

Biographies



Peter M. Adams
BT Networks and Systems

Peter M. Adams is manager of the Mobile Networks and Standards Bureau group in BT Laboratories. He joined what was then the British Post Office as a sponsored student in 1969. In 1973, he became an Executive Engineer working on data standards and was promoted to Head of Group in 1981 on Open Systems Interconnection. In 1991, he became BT's Mobile Standards Coordinator, which included coordination of BT's GSM standards. In addition, in 1994, he took over responsibility for running BT's Standards Bureau.



Nigel Wall
BT Networks and Systems

Nigel Wall is the design authority for the mobile programme within BT's Networks and Systems Division. He joined BT as a sponsored student in 1972, graduating in Applied Physics and Electronics, from the University of Durham, in 1976. During his career in BT, he has worked on satellite systems, telephony, videoconferencing and multimedia customer premises equipment, and, for the past five years, on mobile communications networks and applications. His current focus is on evolution planning for the mobile programme and management of the design of end-to-end solutions involving the convergence of fixed and mobile systems.

The St. Albans Study

The St. Albans study was a practical test to answer two questions. Firstly to see how manual intervention could be reduced in the provision activity within a high-utilisation network design without large-scale re-building. Secondly to determine if the reduction of manual provision activity had a measurable impact on overall fault rate.

Introduction

The BT network in the UK has evolved over many years.

While the regulatory environment and the company have undergone many rapid changes over recent years, the impact on the installed network has been less dramatic and in many situations the network reflects the regulatory environment of a previous era. This is particularly true for the local network from the exchange to the customer.

Significant growth in the network occurred at the time when it was run by the General Post Office (GPO), a part of the Civil Service. At that time, investment in the network was part of the Public Sector Borrowing Requirement (PSBR) and this was a restraining influence on the decisions of the period.

The local network planning rules were affected in two ways. Firstly, cable sizes were kept to a minimum, which, at a time of growing demand, led to repeated incremental build of the network. The installation of larger cables would have met demand for a longer period and therefore needed less frequent additions. Secondly, the network design included many flexibility points to enable the cable pairs to be cross-connected in different configurations

in response to specific customer provision orders.


The general network topology and the associated terms used in this article are shown in Figure 1.

This topology and high level of network intervention in response to customer demand are in stark contrast to a number of other countries, notably the USA where the regulatory environment focused on the return on investment and the cash level of investment was not so constrained.

The parts of the UK where the historic incremental network has not been replaced show high running costs in comparison with more modern topologies. Despite these high costs, the quality of service is acceptable from a customer perspective and the costs are not high enough to warrant the complete replacement of the network in advance of low-cost alternative technologies becoming available.

The high running costs arise from two types of activity, provision and repair. Provision of service regularly requires manual intervention at one of the flexibility points (primary cross-connection points, distribution points, etc.) to utilise available pairs. Quite regularly this requires the removal of a cross-connection between previously used pairs to re-

Figure 1—Network topology



use one of those pairs to serve the new customer. This high level of intervention and disruption of the network creates faults which necessitates repair activity. The repair activity itself creates further network disruption and may lead to further faults. Thus the two activities are linked and together generate the high running costs of this network design.

During a strategic review of the access network, consideration was given to how this situation could be addressed more economically. The review also considered issues associated with the deployment of new technology, which are not covered within this article, but it was concluded that the rate of deployment of new technology was likely to be such that a window of opportunity existed for a copper-based solution to the 'old network' problem if it did not involve complete replacement.

Proposal

The review proposed that the network could be 'stabilised' by further incremental addition of pairs so that provision no longer required physical intervention.

The increment would need to be sufficient to provide for all current unserved properties and make the network design independent of forecast. To enable this concept to be fully tested, a site was selected for a practical study where the network was a traditional design and the penetration was high so that the incremental build would not be too large.

Study description

St. Albans, a small city to the north of London, adjacent to the London Orbital Motorway (M25), was selected as a suitable site to test the proposed solution. The 'St. Albans Study' was initiated in July 1991 to see how interventions to provide service could be reduced and to determine the subsequent effect on fault rates.

St. Albans has a typical mix of business and residential customers with about 38 000 connections.

A study team chosen from the three major Divisions in the company involved in customer provision and repair defined the network build required, specified the processes to be used within a stable network for provision and repair, and specified both new support systems and the changes required to existing ones.

A small group of exchanges was identified with similar network topologies and was left unchanged as a control group.

Network design

Typically, the traditional network design seen at St. Albans had three cross-connection points as shown in Figure 1. The pressurised main cable was provided in minimum 50 pair increments against a forecast of the total demand over the next few years in the area served. Capacity of the distribution cables varied considerably depending on their age, older cables having been provided to a lower forecast assuming a penetration of less than one line per home. Newer cables generally had a higher capacity but were only to be found where recent building had occurred, or where the capacity of older cables had been exceeded.

For the study, all homes without telephones were identified, (approximately 2000) and additional main and distribution cables were installed. This enabled each home to be allocated a pair which was connected to a line card at the switch and terminated either in the home or at the distribution point (DP). The forecast requirement for business customers was allowed for and included in the calculations of the cable sizes, although line cards were not provided in the same way.

The residential pre-connected lines were given a unique class of service (soft dial tone) which only allowed outgoing calls to emergency services and the BT sales office and did not allow any incoming calls.

Where the service was terminated in the customer's home they could simply plug a telephone into the line, call the sales office and be given full service almost immediately by the class of service being changed remotely.

Operational Practices

Advance provision of the network meant that most activities were completed pre-customer order. Thus post-order lead times were considerably reduced and processes very much simplified, with significantly fewer failures. In addition, in the case of residential premises, once the last link had been provided from the distribution point and soft dial tone applied, any subsequent customer movement would be handled without requiring an engineering visit.

Historically, 60–70% of faults occurred in the distribution network. New fault location procedures were introduced to limit intervention and a new decision tree was given to fault repair staff on action to be taken when the fault could not be localised from an initial intervention at the customer's end of the network.

In addition, all soft dial tone circuits were routinely tested and repaired if found faulty to ensure that all circuits were working in advance of potential customer orders.

Finally, to support the changes to fault location practices and limit physical interventions, cabinets containing the primary cross-connection points (PCPs) were locked and access given only on gaining the necessary authority.

Results

Benefits began to be seen from 22 July 1992 as individual PCP areas were completed and the cabinets locked. The network build was completed in February 1993.

The first objective of the study was to reduce the level of manual intervention to provide service. The




Figure 2—Ratio of automated to manual provision

Figure 3—Monthly results





Figure 4—Lifetime fault liability without intervention



measure of success used was the ratio of provision orders requiring no field activity to provision orders that did require manual intervention. The results for St. Albans compared with the control group are shown in Figure 2, and, as can be seen the ratio moved from 1.5 to 5.5 over the period, representing a reduction in field provision costs of 40%.

The increase in the number of orders that could be given a network routing by the sales office also led to savings in planning and design.

By automating so much of the work, the small number of customers' orders not met on the required date was halved, the remainder being due to difficulties on the customers' sites. This has had the further effect of increasing customer satisfaction.

The second objective of the study was to determine the change in fault rate, if any, associated with this reduced level of provision intervention.

The fault rate at St. Albans compared with the control group is shown in Figure 3.

The additional capacity was added to the network in two phases and the fault rate rose considerably during these periods as can be seen in Figure 3. This increase in fault rate during major build can be expected but it is interesting that the increase during the second build phase did not rise above the control group.

On completion, the overall fault rate has been consistently half that of the control group, with a commensurate reduction in costs.

Conclusion

Manual intervention can be reduced by incremental network build and pre-provision to unserved homes. This reduction in intervention does have a significant impact on the overall fault rate.

Neither of these are particularly startling conclusions and could lead to an assumption that this is due to the quality of work associated with provision.

While workmanship clearly does have an impact, the practical observations in St. Albans did not support the view that it is the only factor. The view that we formed can best be explained by considering the whole-life fault liability of plant. It can be assumed that to some degree it should follow the traditional 'bath tub' curve as illustrated in Figure 4.

If this shape does apply across the network, then when a joint or cross-connection point is opened and physically altered to provide service, the whole of that assembly, not just the one pair being worked on, could effectively be moved back to the start of the 'bath tub'. If this is the case with a number of interventions over the element's life, the reliability





Figure 6—Change in fault distribution over PCP areas



curve would change dramatically to a 'saw tooth' shape like that shown in Figure 5.

No practical way of testing this hypothesis was found but if it is a significant factor it would remove the correlation between the age of individual plant items and the fault rate. Because of the incremental nature of the build of the network, it is not possible to isolate plant elements by age but one would expect to see a relatively small proportion of plant (the oldest) giving rise to a disproportionate number of faults. In fact, before the final incremental build, the distribution of faults was over a large number of network elements. Figure 6 shows the distribution of faults over primary cross-connection areas in January 1993, by which time a number of PCPs had been locked but significant work still had to be completed. Data was not collected in this format before that date. After completion of the incremental build and without any

remedial attention to the old plant, the fault distribution changed as shown in Figure 6 by the results for October 1993.

Not only has the total number of faults reduced as referred to earlier but they are less evenly distributed. This suggests the effects of ageing are now a more significant factor and the fault liability of the network has moved from the 'saw tooth' towards the more desirable 'bath tub' shape.

While this study focused on a copper network, it has implications for modern technologies.

Despite the reconfigurability of modern networks, in many cases it is not considered economic to fully equip the initial installation. In addition, the connection to the customer's equipment may be physically separate for each channel or service. In both cases, manual intervention will be required to complete a service provision order.

The challenge will be to account more properly for the maintenance implications of provision in such

Figure 5—Effect of intervention on fault liability

installations and to reassess both network design and equipping policies in the light of those costs.

Biographies



John Prior
BT Networks and
Systems

John Prior joined the British Post Office as a student apprentice in 1966. He graduated from Manchester University, in 1970 in Electrical Engineering, and from City University in 1973. Most of his career has been spent in operations management, although, more recently, he has worked in network statistics, total quality and Breakout. He is Chairman of the Board of Editors of the *Journal* and Vice-Chairman of the IBTE Council.



Kathy Chaplin
BT Networks and
Systems

Kathy Chaplin obtained a B.Sc. (Hons.) in Physics from Sussex University and joined the British Post Office in 1974. After a number of management roles, she has, since 1988, specialised in re-engineering processes impacting the access network. She led the St. Albans study which developed concepts of stability and soft dial tone, and which has resulted in a major roll-out programme within BT. Most recently she has been managing a programme of work concerned with developing a process, which enables the access network to be managed more effectively from access operations units (AOUs).

Developing Sales Engineering into Engineering Consultancy

The sales account-management organisation in BT has been further tuned into line-of-business working. This enables Marketing and Sales to develop business-specific solutions to customers' needs. Within the BT account management teams, the systems engineer (SE) has a major role in developing business solutions that match the evolving technological environment to the equally dynamically-evolving business environment. To this end, the SE is being given the tools and skill-set to undertake engineering consultancy.

Introduction

This article continues and expands a series of articles in the *Journal* on sales and systems engineering^{1,2,3}. The theme is the developing role of the sales systems engineer in shaping and enhancing the application of new technology to solve the challenges facing businesses in the ever expanding and increasingly dynamic competitive environment of the 1990s and into the next century.

The article includes an important contribution from Brian Hillier, telecommunications manager for John Lewis Partnership, one of BT's major clients. As Telecommunications Manager, Brian is responsible for all telecommunications, data and voice, for Waitrose and John Lewis stores (which, as well as retail stores, includes warehousing, logistics, data centres and administrative offices).

We can do no better to end this introduction than by quoting from Brian's contribution.

'The BT account systems engineer has a major role in developing business solutions that match the evolving technological environment to the equally dynamically evolving business environment'.

BT's Sales Structure

Since the previous articles on sales and systems engineering, BT has refined and extended its concept of 'line-of-business working' to include a wider range of client companies and to focus further the account team on business specific solutions.

BT's UK sales force now consists of Global Sales, National Business Communications (NBC) and Personal Communications. NBC consists of Sectors and Volume Sales. The Sectors deal with the largest of the UK client companies, and Volume Sales with small companies. This article concentrates on the systems engineer (SE) in NBC and specifically the SE in the Retail Sales Sector.

NBC Sales, headed by Sales Director John Wheeler, consists of six market sectors:

Retail,
Government,
Industrial,
Commercial,
Energy, and
Finance.

Focusing on the Retail Sector, General Manager Stuart Horwood has over 1000 client companies including some of the biggest in the UK, such as the Post Office, Sainsbury and John Lewis Partnership.

The market sectors are subdivided into lines-of-business manager levels; for example, for Retail:

Transport Services and Distribution;
Retail Stores and Home Shopping;
Food and General Retail;
Retail Leisure;
Retail Travel; and
Retail Post Office and Generic.

Further focus on industry specialism is provided by sub-division at sales manager level, where a sales

Figure 1—A typical sales manager group

manager team of 15–20 sales people will be trained on all aspects of BT customers' business in, for example, Retail Stores and Home Shopping.

The effect of this specific industry focus is that the account teams can develop a detailed knowledge of the industry and its special requirements (for example, specific needs for home shopping businesses to use BT's advanced network services and call centre technology).

To serve these businesses' special needs, a consultancy style of sales team is required to understand the needs and shape solutions to meet those needs. To create this team, the SE becomes a business partner to the account manager in his/her role as consultant to the client, with the SE as technical consultant and the account manager as a business consultant.

Figure 1 shows a typical sales manager team and its relationship to the SE team and line-of-business manager. Although the SEs report to the line-of-business systems engineering manager, they have close relationships with the account team and form virtual teams with the account managers, either on an account basis, or on an opportunity basis, depending on the number and size of accounts managed.


Developing Role of the SE

Traditionally, the entry point for an SE into Sales has been from BT engineering, mainly from a Worldwide Networks background, in transmission or switching maintenance, or from PABX maintenance/installation/project management. This has largely been due to the historical skill base requirement of sales engineering.

Increasingly, the entry point is becoming graduate level with professional qualifications in computing, communications or related technology.

Figure 2 illustrates the pre-1990s and 1990s entry paths for SEs and indicates how SEs can move across to the graduate stream.

Because of the fast developing nature of technologies and service platforms, the customer is increasingly demanding a broader view of



what is available to solve specific and general business needs. For this reason, the SE is developing into a technical consultant with the skills and tools to develop best-fit customer solutions without the initial need to focus design on specific engineering issues or specialist areas until later.


There is a growing trend among system engineers towards internationally acknowledged qualifications. Such routes are Chartered Engineering and professional business accreditation such as MI Management.

The advantages of becoming professionally accredited are manifold; key advantages are:

- internationally recognised as a professional
 - through achievement, and
 - through recognised standards of education, experience and training;
- having access to updated business and technological information; and
- having access to up-to-date concepts in business thinking.

Figure 3 shows an example of how to become a Chartered Engineer.

Figure 2—Systems engineer trends



In addition to the new recruitment profile of system engineers, sales and consultancy training specific to the particular line of business have been developed. This training not only looks at general engineering disciplines but the business skills of an individual and concentrates on developing knowledge of a customer's business/sector. Sector-specific business skills are essential if sector-specific solutions are to be evolved. This training comprises:

- *Product specific training* in new products and product enhancements, specifically to keep the SE at the forefront of new product developments. This training also includes new technology being developed around the world and ultimately involves close links between Sales and BT's research and development at Martlesham Heath Laboratory.
- *Systems engineering development programme* This is a series of training courses to develop the professional skills of the SE. The engineer is introduced further to business planning and strategy.




Figure 3—Institution of Electrical Engineers routes to membership by educational qualification

- **Sales training** General account management and sales training courses to enable the SE to understand fully the sales environment and the rationale behind account strategies.
- **Line-of-business training** Line-of-business training for all SEs (and account managers) delivered by experts in the clients' marketplace, including speakers provided by the clients, to facilitate greater focus on the clients' business needs. Examples include retail marketplace, supply chain management, and inventory management.

Customer Perception

To expand on the developing role of the systems engineer, a summary of observations from Brian Hillier, Telecommunications Manager for John Lewis Partnership, is included:

'The BT account systems engineer has a major role in developing business solutions that match the evolving technological environment to the equally dynamically evolving business environment.

'We perceive the BT SE as our champion, on all engineering issues. Although the account manager has overall responsibility, the technical nature of this environment means that the SE drives many of the issues and solutions through the business to meet our needs.

'We also see the SE as the technical authority for BT solutions, supported by the sales and engineering management, and as the head of the engineering virtual account team for design, installation and service.

'The long-term relationship of the systems engineer to a customer is essential to match the correct solution to business requirements. This can only be done if the systems engineer has a total understanding of our current business and future business directions.

'The way in which BT utilises systems engineers is a clear differentiator for BT when the market place for telecommunications solutions is evolving and growing.'

Developing the Solution

What we would like to do now is walk through two scenarios that develop

following a call from a fictitious client to an account team systems engineer.

'Andy, can you call in tomorrow and talk to me about ACDs?'

The pre-1990s SE would do precisely that, and deliver a technically excellent briefing on the latest developments in automatic call distribution (ACD) and make a fine job of it.

- What location?
- How big does it need to be now/in the future?
- What features do you want now/in the future?
- When do you want it?
- What is the project budget?

These questions investigate the background to the enquiry and dimension it.

The result of this may be an ACD sale.

The consultant SE will also consider why the client has asked about ACD. The SE would follow with further types of questions to identify

the business needs, and discover the implications of the business needs and the impact on the business of fulfilling these needs.

- What does his business need?
- What is the background?
- What business issues are they trying to solve?
- Is there an expectation of further change?
- Which side of the business has it come from?
- Are there any other related business issues that can be solved at the same time?
- How can I add value to the client's business?
- How can I give the client a competitive edge?
- What are the business imperatives and how can I facilitate them?
- What would it cost the business if nothing was done?
- Why is it important to the business?
- What impact would it have on the business if we solved these issues for you?
- Why do callers call?
- What kind of problems are troubling agents and managers?
- What would they recommend be done to make it better?

The result of the above would be a deeper understanding of the client's business needs in this area, and the mapping of business issues to the correctly engineered solution for this prospect and the client's longer-term business. In reality, the consultant engineer would not need to be asked

the initial question about ACDs as he would have been working already with the customer at length on business issues and solutions.

Following from the initial enquiry from our fictitious client, here are two possible scenarios:

Business 1

Business 1 is a retailer whose business consists mainly of retail outlet stores. Ten per cent of the sales turnover is made up by inbound telephone sales in response to sales campaigns. For this reason, the telephone is recognised as a key channel to market. After the visit from the systems engineer, other key points were discussed and found to be pertinent to determine the correct solution:

- Telephone sales constitute a large proportion of business and it is intended that, providing margins remain good compared to a store sale, this business should be expanded. To meet projected growth plans, new marketing initiatives and campaigns are to be planned. To increase effectiveness and improve margins these will need to be focused on the target audience.
- Some of the campaigns will be very seasonal and hence the telephone sales unit must be dynamic to service the response.
- Looking forward 1–2 years some specialist sales will be done over the telephone alongside a greater volume of routine sales. This will mean eventually a lot more than just delivering calls. Control at a very high level will be required to service high-value specialist calls.
- There will be a requirement for regional campaigns and that could stretch into Europe eventually to support the opening of European stores. This will also have the effect of extending opening hours. A 24-hour 7-day week may eventually be in operation.

- A controlled launch is required as the inability to service orders would be counterproductive. All calls are potential sales and customers' high perception of service must be maintained.
- Customer service on the telephone must match high-street store customer service.
- We must make it easier for customers to do business with us regardless of location, therefore it is a business requirement that we extend our channel to market to customers through the telephone.
- A low-risk start up is required that will be truly flexible once we know what the customer response will be.
- A trial launch must be simple and low cost, the response of which will gauge future campaigns.
- To a large extent, there will be an unknown response to some of the new campaigns, especially a new brochure launch. Therefore, it is desirable that risk is jointly taken; that is, initially outsourced for a trial period. The brochure launch is planned for the next quarter to fit in with the marketing business plan.
- We have a reputation of high quality, but we need to monitor this, utilising better management information to maintain and improve our standards.
- We have a limited teleservicing team and we wish to maintain the current site and grow to a second or third site, one of which could be in Europe, without sacrificing flexibility.

Solution—business case 1

It is envisaged that the solution is developed over three stages:

- **STAGE 1**—choose Freefone 0800 sales numbers.

- **STAGE 2**—Use a third-party agency to run a minimum-risk trial for a brochure launch.
- **STAGE 3**—Implement an in-house solution based on collected data.

STAGE 1

On looking at this scenario, there are many possible solutions to meet the business requirements. Several key business drivers point towards a Freefone 0800 sales solution (Figure 4):

- match call to correctly skilled operator;
- maintain and enhance customer service;
- flexible response;
- targeting important;
- growing interest;
- seasonal fluctuations; and
- lots of flexibility required.

By utilising the 0800/0345's advanced services fully, many key business drivers in this application can be realised. The statistics package on 0800 can allow marketing to focus on originating call areas to enable a picture to be built on the level of response from different areas of the country for particular launches. The statistics will also enable problem areas to be detailed to back up the commitment to customer service. Flexibility can be built into an 0800 number delivery plan. Alternative call

treatments can be invoked following peak traffic warnings.

STAGE 2

As the solution develops, it is also apparent that some sort of trial exercise with minimum risk is required.

The company has a brochure that is to be advertised with a new special 0800 number easily associated with the customer and has all of those simple calls answered by an agency with plenty of experience managing large volumes of calls. Connections in Business (CIB) (BT's teleservicing agency) would be ideal, owing to the unpredictable nature of the incoming traffic patterns; CIB is used to working in a very dynamic marketplace and would assist greatly in sizing future applications at minimum risk. CIB can offer economies of scale and resourcing as it is geared to deal with a number of clients' campaigns at the same time. The idea is that CIB would take all advertised enquiries for brochures according to an agreed script and send out agreed information. On a regular basis, CIB would supply full statistics on the operation. Once this trial has been successfully completed, then a first-cut size can be developed, aided by the experience of call-centre specialists, marketing information, CIB information, network statistics, business plan and objectives etc.

CIB will remain in the equation until the in-house call centre is ready to take new traffic. Once the sales team is fully operational, CIB will end its formal trial involvement, but remain available to act as overflow and help in special circumstances such as peak responses to media launch initiatives. (See Figure 5.)

STAGE 3

The platform to support the advanced network services is the next stage. The business drivers indicate that product ranges will diversify, and control of incoming traffic, not only on site but potentially between sites, will become paramount in order to

maintain customer service. This indicates an extra level of processing is required to mirror the business requirements that will change dynamically on an hourly basis and will require different responses, ideally for regional campaigns.

Again there are many ways of looking at this business issue, but for the purposes of the scenario, customer-controlled routing (CCR) based on a Meridian 1 platform with Meridian Mail has been chosen.

BT's Meridian 1 is the integrated services PBX offering in the BT big switch product portfolio. Many applications run on the Meridian 1 platform including integrated mail, automatic call distribution, CCR and integrated voice response.

CCR is an extra layer of processing power that makes the process of call delivery dynamic. This flexible implementation of call routing will benefit customers because call routing parameters can be quickly and easily modified to meet changing business needs. The key to CCR is the flexibility that customers gain because they are provided with control over the call routing decision process (see Figure 6).

Some of the intrinsics that can be used in the call scripts for call routing decision making are:

day of week,
time of day,
number of idle agents,
total number of queued calls,
night service,
calling line identity (CLI) number,

Figure 4—First stage





Figure 5—Economies of scale and resourcing







Figure 6—Customer-controlled routing

day of year,
 age of call in seconds,
 number of agents logged in,
 age of oldest call in queue,
 dialled number information
 service (DNIS) number, and
 calling parties exchange code, etc.

The final solution developed with the customer has been designed as a platform to support all of the identified business needs with a view as far into the future as is feasible. This scenario picked out the basic requirement for a higher level of call processing to match the dynamic market

place that our customer was in. (See Figure 7.) This is by no means the only solution and further investigation would have no doubt brought out some other key considerations. Some other areas that could have been investigated are managed solutions where BT's Communications Management could have managed the call centre, and virtual call centres, where networked Meridians spread over a wide geographical area act as one very large call centre. The CCR part of the solution would enable the incoming call to be answered in the most appropriate way; that is, an

Figure 7—Solution for Business 1




incoming French call would be answered by a French front-end mail message and then directed to a French speaker first. Special campaigns can be answered according to a priority set by the call-centre manager. Service calls can be directed to a completely different call answering script. The calls may be answered by an automatic mail menu or by specially skilled service agents with access to more information. One of the key advantages of CCR is the ability to treat a call individually according to its particular circumstances. Again the possible solutions are endless, but the most important point is that the correct solution can be facilitated only by the correct platform.

Business 2

Business 2 is a retailer and again has a major focus on customer service. This customer has a limited amount of telesales activity, but most of the invited incoming calls are based on customer service activities around the country. The issues faced by Business 2 are slightly different from Business 1:

- Call profiles are seasonal, and history and experience have shown that for this particular scenario our fictitious customer's traffic patterns are predictable and can be staffed appropriately. Contingency planning has up to now coped adequately with traffic peaks.
- Geographically spread customer service desks need individual management. This is a high-cost problem.
- It is costly to provide this service but customer service is a key differentiator in this company's marketplace and consequently must be addressed.
- It is impossible to realise economies of scale that do exist in this company with the current infrastructure.



information, then the option can be given for those calls to be answered by intelligent mail menus. This will allow agents to concentrate on sales and service issues. However, there should always be a safeguard that allows a customer to default to an agent if they require. (See Figure 10.)

To take this one step further might involve using integrated voice response where appropriate, fully automating the process and only involving human intervention where requested.

Another area that could be investigated using this platform is proactive servicing. A list of customers with expiring maintenance contracts could be highlighted each day prompting a call from the service department.

Conclusion

As can be seen, from the same initial question, different business issues have been raised and consequently different solutions have been reached following detailed investigations by the SE.

Figure 10—Use of activity codes

References

- 1 BANKS, COLIN. Account Management in BT. *Br. Telecommun. Eng.*, July 1993, 12, p. 82.
- 2 BANKS, COLIN. Systems Engineering in BT's Account Management. *ibid.*, July 1993, 12, p. 87.
- 3 BANKS, COLIN; and MIKKELSON, MIKE. Systems Engineering in BT's Account Management—The Bid Manager. *ibid.*, Oct. 1993, 12, p. 183.

Biographies



Andy Jeffery
BT National Business
Communications

Andy Jeffery joined BT in 1982 as a Trainee Technician (Apprentice) in London North West District. In 1984, he joined a transmission implementation team and carried on his studies, culminating in an award from BT to study Electrical and Electronic Engineering at University College London. Upon graduating, he obtained a position with BT as a senior systems engineer and joined the NAMGRAD (National Account Management Graduate Programme) training programme in 1991. After 4 years of structured and pre-defined education, training and experience, he has been elected as a Chartered Engineer and continues his work as an account-based consultant systems engineer within Business Sales.



Colin Banks

Colin Banks joined the then British Post Office in 1963 as a Trainee Technician (Apprentice) in London North Telephone Area. In 1976, he joined the Strowger local exchange maintenance division and progressed to managing local Strowger and TXE4 exchanges. In 1981, he became responsible for the Area's Management Services Unit, including manpower planning for the local exchange modernisation programme, and area computing services. In 1984, he became part of the team that built National Networks (latterly to evolve into Worldwide Networks), and ran transmission and trunk switching operations for Northern London District until 1988. He then moved into BT national account management as a senior project manager for the Lloyds Bank account team, and in 1991 became systems engineering manager, initially for the NatWest Bank account team and then for the London-based banking line of business. In 1993, he moved within national account management to be the senior bid manager for the global banking line of business, progressing in 1994 to run the NBC Retail Sales sector bid management unit, before leaving BT in March 1995.

Telecommunications Fraud Management

Reducing Costs and Winning Business

With competition and regulation in the European telecommunications market biting harder each year, suppliers need to establish new ways to keep ahead in terms of cost control and market differentiation. Many telcos now take fraud seriously. Those who do not will see their costs rise and their customers leave. This article briefly outlines the history of telecommunications fraud and the evolution, possibly revolution, necessary to manage fraud in the future, helping to secure market positioning for those who heed the signs.

Introduction

Fraud can be seen everywhere. Finance houses, government departments and other major business institutions often make the headlines through their exposure to fraud. Telecommunications suppliers are particularly vulnerable. With their race to win customers from the competition, the warning signs can be missed or ignored.

Does it matter? With fraud losses representing sizeable portions of new customer revenues, profit margins are further eroded and bad debt levels increase. Confidence in the supplier is greatly reduced if a customer is hit by fraud. Ultimately, major customers will, and should, migrate to the supplier who can best protect them at reasonable cost.

What is Fraud?

Fraud can legally be defined as 'A dishonest act that prejudices another person's (economic) right'—causing intentional loss to another by dishonest means.

Primarily, the fraudster intends to make money. This can be achieved in many different ways—selling calls on cheaply (usually to international destinations), selling codes that allow 'free' calls to be made, utilising services that have a cash benefit associated (for example, premium rate services (PRSs)) and so on.

It may be that a service is obtained purely for fraudulent use. This itself may vary from a simple telephone line to calling cards (Chargecards) to business switching apparatus to ISDN to..... and so on.

An existing customer, product or service may also be targeted. Many fraudsters will attempt to duplicate or steal a customer identity to gain access to his or her service and make free calls. Examples of this have been seen in the analogue mobile networks, calling card services, business switches, private or virtual private networks; in fact, anywhere an access code or customer identity is utilised. Potentially, these types of fraud are the most damaging as genuine customer accounts are affected.

In fact, most telecommunications products and services present an opportunity for fraudulent use, an opportunity to make free calls, an opportunity to make money.

The common fraud factor is the intent not to pay for the call or the associated service or product being used. This is **the** distinction between fraud and bad debt—if there is intent to avoid payment while using the product or service, then this is fraud. If the intent to avoid payment arose after usage, perhaps because of an unexpected lack of funds, then this is simply bad debt. This distinction can be difficult to make but with a few clear guidelines and criteria to watch out for, it can usually be made.

The Impact

Frauds directed at existing customers could destroy them (financially) or cause massive losses to suppliers who are prepared to take the brunt of such attacks. Either way, confidence in the telecommunications services used will be impaired in the eyes of the customer affected. They may well leave and could be major revenue earners.

Historically, suppliers of telecommunications products and services have been the primary target for fraud. Defences within the industry have changed to reflect this focus. In many cases, particularly in the larger organisations, a lack of coordination and cooperation has simply led to a shift in the fraud type or method. Fraudsters will migrate to wherever it is easiest to thrive. Telcos are a type of 'culture' for the fraud 'bacteria'—if only one area is cleansed, the bacteria simply grow elsewhere.


The effect of ill-conceived ideas or the unilateral approach of some business units within a telco may well be to disperse fraud, perhaps even make it invisible. Figure 1 highlights the possible effect of such activity.

Although Figure 1 may not accurately reflect the proportions of visible versus hidden fraud, the iceberg theory is perfectly legitimate. Telcos do not usually know or measure all fraud impacting upon them or their customers. The proportions would depend on several factors:

- systems used to monitor fraud levels,
- reporting and measurement mechanisms used,
- the degree of training and awareness given to company employees, and
- the desire and motivation to make fraud levels visible.

Measurement is discussed in more detail later.

Figure 1—The possible effect of ill conceived or unilateral fraud controls



Costs

Telecommunications fraud within the USA during 1993 is estimated to have cost \$3–5 billion. In the UK, total losses due to fraud are kept in close commercial confidence. However, losses are significant and run to millions of pounds each year. Similarly, the European market suffers at the hands of fraudsters, but to what extent? Again, fraud levels are kept secret but European fraud is significant. This will grow; as former communist countries evolve, so will their telecommunications market and with it their element of fraud. This is inevitable but it can be managed.

Revenue loss is not the only cost to consider. Others include:

- cash payments to PRS providers, interconnect payments to other licensed operators (OLOs) and international settlements;
- cost of high-value telephony equipment lost through fraudulent applications for service;
- resource spent in processing orders, engineering provision and maintenance of service, credit, risk and debt management activity related to fraudulent accounts;
- cost of poor network and business planning/strategy decisions based on high levels of fraudulent traffic (expensive international gateway switching may be installed, inadvertently facilitating fraud); and
- the cost of lost customers though a breakdown in supplier confidence.

In most cases these can be quantified. Measurement is a key element of any effective fraud management strategy.

Fraud Measurement

Why is it necessary to distinguish fraud—surely debt is debt? Fraud

must be identified separately so that it can be tackled effectively; that is, by root cause analysis. To do this, fraud must be identified by type, method, product, service and customer and then measured accurately and consistently. Only by making it visible will it become possible to address. Of course, once it is visible, it is possible to then measure how effective countermeasures have been. It will help business units make the right marketing and strategy decisions.

Measurement should be considered in three parts (I³):

- **Identification** Through training and heightened awareness, along with access to appropriate tools, levels of fraud recognition will be improved within the organisation.
- **Information** Information from the frauds identified should be captured in sufficient detail so that it can be of use in later analysis. Within the measurement mechanism, it is important to maintain structure and flexibility in combination with simplicity of use and application. The process may be facilitated by direct feeds from fraud detection and management systems.
- **Inclination** Motivation is fundamental. Product and business unit managers should be committed to fraud measurement to make it work. Reluctance to identify and disclose fraud information on the basis of perceived poor performance needs to be countered in order to obtain an accurate picture. Methods of measurement are also important, for instance measuring actual rather than billed revenue as the measure of the worth of products and services.

If the measurement process is cumbersome, people will not use it, and statistics will be affected. Automation and simplicity are key in

As with any business, stifle the fraudsters revenue opportunities or increase their costs, and they will soon feel the pull of liquidation.

this instance. Paper-based systems work but take time and effort. A simple key-stroke when a fraud is identified is far more likely to be used and accuracy is enhanced.

Measurement is the first step along the route to a successful strategy. The focus must now turn towards what needs to be done to reduce the effects of fraud.

Fraud Management— Current Situation

There are many ideas on how to control fraud. These range from simple credit-vetting of customer applications to computer-based call/traffic monitoring systems. In essence, none of the methods currently used are particularly sophisticated. Even in the USA, where fraud detection has been evolving for some years, the systems are fairly simple. Many would argue that simplicity is the key. The trouble with simplicity is that it often correlates to inflexibility. Fraud is continually changing and flexible systems need to be employed to combat it.

Before the availability of current technologies, the primary means of defence against fraud was by means of post-event investigation. There were some great successes. Although seen by many as an effective cost-control strategy, in reality the cost of manpower and the continuation of fraudulent accounts for detection purposes would result in high costs being incurred. Also, this did little to deter the major criminals—telecommunications crimes attracting lower sentences and less risk for similar return than alternative crimes. Furthermore, the removal of one or more criminals would leave a swiftly and easily filled vacuum in many cases.

With technologies abounding, remote and easy network access techniques make it more difficult to manage fraud manually. Furthermore, legislation has not kept up with technological revolution, and becomes an investigative constraint. Those glory years may be dwindling.

Businesses need business-like action. Delayed service termination, high resource costs and uncertainty of outcome make it difficult to rely solely on investigative action. A better deterrent is to hinder significantly the money-making potential of a fraud so that it is no longer a worthwhile project. As with any business, stifle the fraudsters revenue opportunities or increase their costs, and they will soon feel the pull of liquidation.

This is where fraud detection systems come in. Monitoring of access, billing or call record information is the commonest method—looking for high or unusual usage (or attempted usage). The quicker this can be done, the quicker fraud can be detected. Current systems typically review billing data daily, with some achieving checks every few hours. Ideally, however, real-time analysis and alarm generation must be sought, while at the same time reducing the likelihood of false alarm generation and action.

If fraudsters were consistently hit within a couple of hours of operation, they are unlikely to continue their attack in the way they had been, and possibly not within the telecommunications industry. But there are dangers with simple monitoring systems. Whatever parameters are built in, they may soon be recognised and skirted. That is why systems need to be flexible and driven by continuous analysis of results.

Fraud monitoring should not be the only focus of fraud control. Greater emphasis should be placed on prevention of the fraud in the first place. While this may seem obvious, it is continually overlooked. Products, processes, systems, operations can all be defended in some way, usually for little or no cost, particularly if incorporated at the initial stages of specification. Customers will be pleased to hear that it is secure too—marketing opportunities abound.

- Building in security costs less than bolting it on later.

- Defences are often far more robust if they develop with the product or service.

Current fraud problems are usually accentuated by a lack of cooperation and/or coordination within the business. Defences of one form or another are often developed without considering the ripple effect. The best that can happen is that the problem is passed to a competitor. Usually, however, it is moved within the business to another department, product or service.

Competitive displacement of fraud may be seen as an attractive strategy. It may be.... in the short term. But it should be borne in mind that telco people and ideas move around the industry as well as the fraudsters. There may be a recurrence, which could be even more damaging because of the changing methods of attack used to avoid the highly developed defences. Avoiding the see-saw is an attractive alternative. Ridding the **industry** of high-cost fraud problems will benefit all, the industry's customers in particular.

Fraud Management Strategy



Key elements of an effective fraud management strategy are:

- *Policy and guidance* The company must have the framework on which to build.
- *Measurement* Knowledge about where the problem is and its size is essential to determine how the problem is to be fixed. Quantifying the success of the fix is also necessary—it may need some more work.
- *Training and awareness* Almost anyone within the industry has the potential to spot a fraud in action, or a weakness which could lead to fraud. Raising the profile and giving them the tools to highlight problems is essential to continuous improvement.

- **Ownership, accountability and co-ordination** These are difficult to implement in reality. This is primarily a cultural issue. People will say one thing but do another. Linking improvement to performance of loss owners is the best way to ensure cooperation. A single, high level, point of ownership is also important.
- **Layered defences** No organisation should put all of its eggs into one basket. The 'peeling an onion' approach to physical security is valid in the fraud management environment too. Defences should be present at the front end, within the operational aspects and as a back stop firmly propping up the above. See Figure 2. Billing and collections processes are also prime for fraud management development.
- **Security in product development** Defences must be built in at the start. Use should be made of the expertise available—it will be costly to ignore. Many initiatives have marketing benefits and a good consultant will balance business requirements against threat. Security need not be a barrier. It can be a differentiator.

There are other elements to consider depending upon the nature of the organisation; plan ahead, agree the strategy with all of those involved and review it regularly. If implementation of certain elements is not working, it may be necessary to change course. Frequent change should be avoided,

Figure 2—Layers of security application



but developments allowed to determine what should be altered. If any weaknesses are left, it can be guaranteed they will be exploited.

Investigation may well form part of the strategy depending upon the other elements present. If the systems and improvement cycle are swift and robust, there is a reduced need for investigative resource. Even so that resource could be used to target fraud organisers, the kingpins who rake in fraud profits. The trouble is, these are often based abroad and international laws are even more difficult to negotiate than in the UK. There is no doubt that investigation is costly but it does have its benefits (targeting, crime intelligence, *modus operandi* etc.). These need to be balanced within a strategy.

One of the main defences should be a *fraud management system* (FMS). This is not a simple detection tool, but the next generation of detection tools coupled with a host of management capabilities. The focus is to manage the problem out, not simply detect and deal with individual frauds as they occur. The simple 'fire engine' approach is not valid for the future of fraud. FMS aspects will also save on resource where simpler detection systems require a high degree of manual intervention. Customers too will expect a degree of protection. The better the service, the greater the chances of retention and revenue gain.

New technologies and world marketing strategies open up the UK

and European telecommunications industry wider than ever before. There will undoubtedly be a shift in fraud patterns and methods reacting to the new opportunities and defences deployed. The penny scams of today might be the million-pound losses of tomorrow. The new generation of systems need to be able to cope with this change—continuous change.

Not only will the FMS detect and manage individual frauds, it will link apparently unassociated frauds or customers, profile fraudsters and spot them when they attempt to gain access or incur usage. It will provide management information to take action and minimise losses in areas where fraud was not previously suspected. It will intelligently 'roam' the network proactively seeking out fraud. It might be seen as a network guard-dog.

Continuous analysis of results will provide for new fraud detection parameters. New systems will suggest changes to the defences employed in the layers above—perhaps a different question set for new customers, perhaps new usage limits for calling card holders and more. As with the automated fault finding and correction capabilities of new switching apparatus, the move will be towards automated fraud management as operational confidence builds. Figure 3 gives some examples of such activity.

The FMS is very much an improvement tool—it should not be seen as just a fraud weapon. For example,




Figure 3—Possible input and output streams of a fraud management system

there is significant overlap with the credit and debt management environments. Why build many different systems to perform similar tasks? After all, monitoring is the basic task in all these areas. Intelligent interpretation of data is no longer the bastion of the human being. These business areas have useful skills for the fraud arena. The activity overlap is significant and cannot be ignored—some frauds will not be as blatant as others. Many debtors might later appear under false names. It makes sense to work together, to focus resource and efforts and thus improve performance.

Conclusions

Fraud of one form or another will represent a significant portion of a telco's bad debt. Whether people within the organisation accept that or not, it's a fact. It may simply be unrecognised or worse, hidden.

Good customers too are affected by fraudsters. Perhaps fraudsters may disguise themselves as corporate customer employees and piggy-back customers' services. Perhaps an employee will sell on calls or access codes to others for cash. Telcos cannot afford to miss the opportunity to help manage their problems. Telcos may be forced to accept the losses in any case, if not through goodwill, then perhaps through litigation.

A plethora of opportunities present themselves. Telcos can reduce their costs and improve their ability to win business, generating new revenue. But competition should not be at the expense of the future, which may well

hold higher losses for the organisation tending towards competitive displacement.

A key factor is to communicate—with those who have the expertise and experience of how to manage the problem out, possibly with competitors and, most importantly, with others within the company. A united approach can work. Customers will come to expect it to work.

Biographies



James Messham
Fraud Management
Limited

James Messham, a fraud management consultant, is currently studying for a combined bachelors/masters degree in Electronic Communication and Computer Engineering at Bradford University. He has undertaken several placements with BT in technically related projects making best use of his specific learning and skills. In particular, he was recently responsible for the development of a software-based modelling system providing a fraud sensitising capability to telecommunications business managers.



Phil Kelly
Fraud Management
Limited

Phil Kelly, a fraud management consultant, finished university with a degree in Biochemical Sciences and then joined the police. He joined the Post Office Investigation Department (POID), which resourced inquiries into telecommunications fraud in the

mid-1980s. In 1989, Phil moved to BT's Investigation Department and has since investigated most types of telephony crime. After focusing on serious organised crime, he was appointed Crime Intelligence and Analysis Manager where he introduced a variety of software solutions. He completed an MBA with the Open University before joining the Commercial Security Unit of BT where he has contributed to the specification of fraud detection and management systems and business procedural security reviews.



Dean Smith
Fraud Management
Limited

Dean Smith, a fraud management consultant, joined BT Investigation Department in 1989 from HM Customs and Excise Collection Investigation Unit, Birmingham. After two and a half years managing investigations into a variety of telecommunications-related crimes, he joined BT's Commercial Security Unit specialising in fraud prevention and management. Dean's early efforts concentrated on measuring the fraud problem and building a rigid framework for future fraud management development. This included the design and implementation of a fraud identification and measurement training package and advising on the development of fraud control mechanisms in the business. Later work includes the invention of the debt reduction system, a call 'fingerprinting' technique to spot fraud and debtors, and secondment to the Breakout programme building a fraud management strategy for BT. Latterly, Dean moved to Mercury Communications Limited in a similar capacity.

High-Voltage Protection of Telecommunication Circuits

Some electricity generating and transformer stations can be subject to an excessive rise of earth potential owing to faults on the associated power lines. Copper pair cables which serve these sites may present a hazard because they introduce a remote earth via the telephone exchange. This article describes the precautions necessary when providing telecommunication services to such sites.

Introduction

Electricity generating and transformer stations can experience a local rise of earth potential during certain types of fault on the electricity supply network. At some sites, particularly those in areas of high soil resistivity, the rise in potential can reach several thousand volts relative to a remote earth.

The electricity companies require telecommunication services to their sub-stations for telephony, telemetry and the control circuits of their power switching and protective systems. They are also interested in using ISDN2 for remote video surveillance.

If these services are provided over copper pair cables, there is a danger to both personnel and equipment during a rise of earth potential due to the large voltage difference that may exist between the exchange end and the customer end of the circuits. Telecommunication network

operators are therefore obliged to adopt protective measures on lines to sub-stations which are subject to a large rise in earth potential. Figure 1 shows a typical site where high-voltage protection is necessary for telecommunication circuits.

At some sites the private properties adjacent to the sub-station are also affected by the rise of earth potential. Telephone lines to these properties may require protection.

It is in BT's interest to be able to offer services to sub-stations, even where special protective measures are required, because, taken together, the electricity companies form one of BT's largest customers and it is very worthwhile keeping their business.

This article explains how rise of earth potential occurs and describes the measures taken on telecommunication lines to ensure the safety of people and equipment.

Telecommunication lines to a sub-station may run parallel to the power lines into the site and can be

Figure 1—A typical electricity generating station
(Photograph courtesy of the Electricity Association)



subject to voltage surges induced by the large out-of-balance current flowing in the power lines during a fault. Nearby lightning strikes can also cause high-voltage surges to be induced into telecommunication lines. However, that is another story and this article will confine itself to the problem of rise of earth potential at the sub-station.

Rise of Earth Potential

At sub-stations a low resistance connection with the general mass of the earth must be installed. This is to provide a return path for fault currents from the power lines and for protection purposes. It is sometimes difficult to obtain a low resistance connection with the general mass of the earth, particularly in areas of high resistivity. When a substantial earth fault current flows, it is possible for the metalwork which forms the earthing system to acquire a high voltage relative to a remote earth. It will exist for the duration of the fault, up to one second, and is equal to the product of the earth fault current and the impedance of the station earthing system through which it flows.

The rise of potential of the earth system during a power line fault does not normally endanger the personnel in the station since it is general practice to bond all the non-current carrying metalwork within a station to the earth system. Thus all the station metalwork is at the same potential and personnel cannot come into contact with two sets of metalwork at different potentials. However, telecommunication circuits entering into sub-stations are usually earthed at the telephone exchanges and so their conductors should therefore be considered to be at true earth potential.

If a heavy current flows through the earthing system at the station the consequent rise in earth potential can cause two hazards:

- the risk of electric shock to a telephone user in contact with

metal work connected to the station earth system, and

- the risk of damage to the equipment connected to the incoming telephone cables.

Furthermore, should telecommunication circuits require a connection to earth at a sub-station for operational reasons and this connection is near the station earth system, then, under fault conditions, its potential may rise with consequent risk of shock to personnel and damage to telecommunications plant remote from the station.

Telephone lines provided for customers adjacent to sub-stations may also be subject to hazardous voltages, if the equipment itself or a person using it is in contact with metalwork at local earth potential; for example, radiators and water mains. This can continue for the duration of an electricity supply fault current flow.

The use of optical-fibre cables eliminates many of the problems associated with rise of earth potential. BT has developed a type of optical fibre cable for use at hot sites which has no metallic parts; this can be compared with standard optical-fibre cable which has a metallic

strength member and possibly a metallic moisture barrier.


MegaStream services to hot sites are provided over optical-fibre cable. Other types of service also use fibre cable where the volume of circuits justifies it. At present it is not economical to provide a small number of circuits over a fibre cable.

Assessment of Rise of Earth Potential

Currents flowing in the ground as a result of a phase-to-earth contact in a high voltage power line will cause a rise in earth potential, relative to a remote earth in the regions where current enters and leaves the ground. To enable the precautions for safeguarding telecommunications plant against rise of earth potential to be determined, it is necessary to assess the magnitude of the potential rise of the station earth system under fault conditions. Figure 2 shows, in simplified form, the conditions which arise when a fault occurs on a system whose neutral is earthed at the station.

An earth fault on a line fed from an electricity station will cause a line fault current I_1 to flow. This current I_1 will flow through the earth unless the power line is provided

Figure 2—Fault conditions on a system whose neutral is earthed



with a continuous earth wire or a metallic cable sheath bonded to the earth electrode system. In which case, a return current I_2 will flow in that path and the current through the earth will be $(I_1 - I_2)$. The currents $(I_1 - I_2)$ and I_2 flow back to the system neutral point over the paths shown. It is the earth current $(I_1 - I_2)$ flowing through the impedance Z of the station earth system which raises the potential of the system relative to that of the true earth. The rise of earth potential will thus be $(I_1 - I_2)Z$ volts.

For a 400 kV power line, the total fault current can be as high as 60 kA, with 20–30 kA flowing through the ground.


When the fault is at a remote station (not shown on Figure 2) the earth electrode current will produce a potential rise in the earthing system at both stations.

If a fault occurs within the station itself the result is different. These conditions are shown, in a very simplified form in Figure 3. The fault current I_3 , fed from the local transformer(s), does not produce a rise in earth potential as it flows through the station metalwork to the system neutral point.

Soil Resistivity

Soil resistivity is the resistance of soil and the most common unit of

Figure 3—Fault conditions when fault occurs within a sub-station



measurement is the ohm-metre. This is defined as the resistance across opposite faces of a one metre cube of earth. Typical values range from marsh lands 2–3 ohm-metres; sand 90–800 ohm-metres to rock 500–10 000 ohm-metres. The levels of strata under the surface will produce significantly different values of resistivity and these values will change with both the amount of moisture in the soil as well as its temperature. The resistivity of frozen soil is particularly high. It is because of the substantial resistance of the various types of soil, and particularly rock, that large currents through them produce a considerable rise of earth potential.

The map shown in Figure 4, on the next page, shows the wide variation of soil resistivity in the UK. The higher levels occur in the more mountainous regions.

This map takes into account the levels of strata down to 150 m.

Levels of Rise of Earth Potential to be Expected

The rise of earth potential is the difference between the potential at an electricity station during an earth fault on a power line and the potential of 'true' earth at the remote end of the telecommunication cable. This rise in potential is produced by the actual fault current flowing through the station's ground grid impedance to remote earth. This current will not be the total fault current used for the design of circuit breaker capacity.

In areas of high earth resistivity, the rise in earth potential may reach 10 kV RMS during an earth fault on a power line.

Acceptable Limits of Rise of Earth Potential

Special precautions against rise in earth potential are normally only required in respect of sub-stations which are classified as *hot*. A hot site is one where the rise of station

potential exceeds 430 V RMS under the most severe fault conditions. The limit of a 430 V RMS may be extended to 650 V if the electricity company confirms that all the power lines contributing to the earth current, from which the rise of earth potential is assessed, are in the 'high-reliability' category. High-reliability lines are those operating at 33 kV or greater voltages and are controlled by switchgear provided with main protection equipment which will disconnect a fault current in less than 500 ms, and usually within 200 ms.

'Hot' Sites

Areas within which the earth potential may rise above prescribed limits described above are known as *hot zones* and the site of the sub-station is known as a *hot site*. Any part of a zone which is outside the boundary of the electricity supply site is called *off-site hot*. The rise of earth potential falls away as the distance from the station earthing system increases.

The hot zone is defined as the area around the electricity station bounded by the voltage contour line where the voltage reaches either 430 V AC or 650 V AC as appropriate. Traditionally, in the UK, the boundary of a hot zone has been taken as a 100 m around the perimeter of the high voltage compound (Figure 5). However, the radius of the hot zone can be established with more accuracy by calculation or by current injection tests. It is worth determining the true radius of the zone if this is less than 100 m and in consequence leads to a reduction in the amount of protection required. In some cases the radius of the hot zone is much greater than 100 m. For example, at the Dinorwic pumped storage station in North Wales, the average radius of the 650 V contour line is about 1 km.

Earthing at Hot Sites

At all sub-stations an earthing grid is installed to provide a common




Figure 4—Map of soil resistivity in the UK




Figure 5—A 'hot zone' representation

telecommunication cables to ensure security.

Isolating links

At generating stations and sub-stations operating at 275 kV or 400 kV, and at all hot sites, telecommunication cables must always be terminated on isolating links. A drawing of Link Isolating 6E is shown in Figure 6.

At non-hot sub-stations operating at or below 132 kV, the termination of cables on isolating links is not normally required in areas of low earth resistivity. However, in areas of high resistivity, such as in North Wales, isolation links are always provided. The isolating link enables the incoming circuit to be disconnected while service personnel work on the cable or the line isolating unit.

Wiring from the isolating links to the BT line isolating unit (line side) is made by individual 2-wire cables (0.4 mm diameter). After the line isolating unit (customer side), the cabling is always at local earth potential and thus does not need special insulation. However, this cable should not run against or

ground for all electrical equipment and metalwork in the station. The earthing grid consists of bare copper conductors buried 300 mm to 500 mm below the surface in a grid pattern. At cross-connections the conductors are securely bonded together. In addition, earth rods may be installed at the corners and at junction points along the perimeter. The metal cases of all line isolation equipment must be bonded to the station's earthing system. No connection is permitted to be made from the line side of isolation equipment to the station earth.

within a hot zone from working on cable pairs that are at the remote exchange earth potential.

Overhead lead-ins are not permitted at electricity stations operating at 132 kV or above. On sites operating below 132 kV, lead-ins are always, where possible, carried in underground cables to ensure the highest level of reliability. An electricity station may require two or more physically separated


Protective Measures on Circuits Serving Hot Sites

A combination of measures is taken on telecommunication circuits serving hot sites to secure the safety of BT's customers and its network personnel, and to protect equipment connected to the circuits.

Cables entering hot sites

Special rules are applied to cables serving electricity stations. Underground cables with metallic pairs must have copper conductors but must not have a metallic moisture barrier and no joints are allowed within the 100 m hot zone. This requirement is to prevent people

Figure 6—Link Isolating 6E



across the incoming cables which are at remote earth potential.

Isolation equipment

At all sites which are classified as hot, equipment is provided to electrically isolate wiring at the station earth potential (customer side) from the incoming telecommunication cable pairs which are at the remote earth potential (line side). The isolation equipment is located as closely as practical to the point at which the isolating links are located. There is a range of line isolating units for the different types of service. In view of the importance of maintaining continuity of service, these units are designed to ensure that the circuit and equipment will function normally both during an earth fault on the power lines and after the fault current has ceased.

The BT requirement for line isolating units is that they are rated to withstand 10 kV AC applied longitudinally between the line side and customer side windings. They are tested to withstand 20 kV RMS at 50 Hz applied for one minute between the line side terminals and the customer terminals which are connected to a local earth.

Two methods of isolation are used; a transformer and a short optical link. The type of transformer used is sufficiently well-insulated to withstand a high common mode voltage applied between the line (exchange) side and customer side windings. The transformer carries the AC signals (voice or data), any line signalling or supervisory conditions being transferred by optical means.

In the other method, line side electronic circuitry is linked to the customer side circuitry by a number of optical fibres about 10 cm long. All signals are coded/decoded by the electronic circuitry and transmitted over the fibres. Isolation is provided by the non-conducting gap between the line side and the customer-side circuitry.

BT currently uses the line isolating unit (LIU) 3B for exchange lines and the LIU 8A for analogue private circuits. The LIU 3B uses the optical-fibre link method of isolation. The LIU 8A uses transformers and can handle audio-frequency signals only.

With the advent of a new generation of KiloStream equipment and ISDN2 which use different line codes, it was decided that it would be desirable to have one unit to deal with both requirements. It was thought that it would be easier to develop a single transformer based unit, rather than a single unit having electronic circuitry capable of dealing with several different line codes. However, it proved more difficult than expected to develop this unit, mainly because the additional components required in the case of ISDN2 to provide line power to the network terminating equipment affects the transmission properties of the transformers. As an expedient, the Line Isolating Unit 9A was introduced for KiloStream only while development continued on a unit suitable for ISDN2.

Exchange lines—Line Isolating Unit 3B

This unit may be installed in an electricity station or on customers' premises which lie within the hot zone of an electricity station. This

unit is suitable for providing isolation for the following customer services:


- direct exchange lines (loop-disconnect or multifrequency signalling),
- payphone lines (Prepayment Only),
- PBX lines,
- PBX extensions, and
- subscriber private metering (SPM).

The Line Isolating Unit 3B, shown in Figure 7, is located within a maximum distance of 1 km of 0.5 mm cable from the customer's equipment. The inside of the unit is shown in Figure 8 where the fibre-optic couplers can be seen. There is provision for setting the facilities

Figure 7—Line Isolating Unit 3B



Figure 8—Inside Line Isolating Unit 3B






Figure 9—Line Isolating Unit 9A

required which are preset and no adjustment is required on site. Apparatus that is approved for connection to the network to provide any of the services listed above can be connected to the customer side of this line isolating unit and the customer will have all the facilities normally available with the apparatus.

The line-side circuitry in the isolating unit is provided with power from the exchange line. The customer's side of the unit is powered from a local 50 V dc supply or 240 V AC mains or both. The unit will provide ringing with either type of supply; but mains must be provided for payphones and the SPM service.

KiloStream Data Circuits—Line Isolating Unit 9A

KiloStream is a range of private digital data circuits offering network synchronised duplex transmission at speeds of from 2.4 kbit/s to 64 kbit/s. It is designed for high-speed, high volume data transmission and, when it is sub-multiplexed at lower kilobit/s rates, it can carry a mix of speech and data applications in a single data path. It also provides the transmission link for high-speed fax and slow-scan TV.

KiloStream circuits can be protected by Line Isolating Unit 9A

(Figure 9) connected at the substation between the incoming line and the network termination point (NTP). In order to maintain safety to personnel in the substation the installation must be carried out in accordance with specified rules.

The line isolating unit comprises two discrete isolation transformers. These transformers are manufactured to a very high degree of accuracy to provide minimum loss at frequencies up to 200 kHz. The unit provides isolation of up to 65 kV peak and 20 kV RMS continuous while maintaining full communication at up to 56 kbit/s for either a 4-wire or two 2-wire data circuits. The unit does not require power from either the exchange nor from an external supply.

The line isolating unit (Figure 10) is also available as a plug-in card, which can be inserted in an eight-card shelf. The advantage of the eight-card shelf is that up to eight 4-wire circuits or sixteen 2-wire circuits can be protected by one shelf unit.


ISDN2 Circuits

The BT ISDN2 service is available to customers served by digital exchanges and uses a 2-wire connection into the local network, identical

to that of an ordinary PSTN line. The maximum reach of a BT ISDN 2 circuit is similar to that of a PSTN line. Once within the digital exchange, an ISDN call is switched and transmitted in the same manner as any other call. This service delivers an information rate of 144 kbit/s which is broken down into two 64 kbit/s channels for customer communications (also known as 'B' channels) and a 16 kbit/s channel for signalling on the network (also known as a 'D' channel). The increased bandwidth means that non-voice applications which are beyond the bounds of the switched analogue environment are now available; for example, group IV facsimile, CAD/CAM, file transfer, picture transfer and video-conferencing.

A line isolating unit is being developed for the protection of ISDN2 circuits which will be similar to that for KiloStream circuits. However, because of the signalling requirements of the ISDN service, it is necessary to provide a fibre-optic link for signalling in addition to a high-voltage isolation transformer for the data signal. The exchange-side circuitry is powered from the exchange while a local DC source provides power for the station-side circuitry and feeds the station-side line with a current of up to 50 mA

Figure 10—Line Isolating Unit 9A in plug-in card form



at a voltage of up to 150 V for the customer's apparatus. As with the KiloStream line isolation unit, the ISDN line isolation unit will provide protection against 20 kV RMS and 65 kV peak.

The Future

It is already standard practice to use fibre cable to provide a MegaStream service to a hot site. As fibre achieves a greater penetration in the BT network it should be possible to provide all services to hot sites over fibre, thus avoiding the need for the special precautions required with copper pairs cables.

Conclusion

Excessive rise of earth potential at electricity stations poses a serious threat to the safety of personnel and equipment. The practices adopted by BT have successfully prevented injury and damage over many decades of providing services to hot sites.

Acknowledgements

The National Grid Company for permission to photograph the line isolating units in a sub-station. Positron Industries Inc. of Montreal, Canada for their assistance in the production of this article.

Biographies



John Gerrard
BT Networks and
Systems

John Gerrard joined the GPO in 1964 as an Open Competition Assistant Executive Engineer. He subsequently obtained an honours degree in physics at Birmingham University. On his return from university he was appointed to a post in Telecommunications Headquarters where he worked on the maintenance of coaxial line systems. After promotion to Executive Engineer in 1979, he took up a post in trunk transmission planning. More recently, he has been working in the area of electrical protection. This includes all types of electrical problem from low-voltage interference to avoiding hazards to BT people who are required to work in the vicinity of high-voltage plant.



Walter Sweetenham
Positron Industries
Inc.

Walter Sweetenham graduated in Electrical Engineering at Bristol University and then joined the Apparatus Division of Standard Telephones and Cables at their New Southgate factory. He later joined the Railway Signalling Division of Westinghouse Brake and Signal Co Ltd. at Chippenham. He is now the UK representative of Positron Industries Inc. of Montreal, Canada. Walter is a member of the Institution of Electrical Engineers and a Fellow of the Institution of Railway Signal Engineers.

The Development of Power and Building Engineering Services in Telecommunications

The adage 'no power, no service, no revenue' is as true today as it was 50 years ago. This article reviews the development of telecommunications power and building engineering services since the 1930s and discusses future requirements.

Introduction

Telecommunications power and building engineering services (BES) plant has come a long way since the 1930s.

On the DC side, equipment has developed from on-site generation of DC power to distributed switch-mode power supplies. AC stand-by plant has progressed from zero provision to fully automated packaged engine sets, and ventilation requirements have advanced from openable windows to temperature controlled sites with microprocessor controlled air handling units (AHUs).

Advancements in technology, energy considerations and careful design have all played a part in producing AC, DC and BES plant that is purpose built to support modern switches reliably.

In recent years there has been continual reduction in available resource and this, coupled with the privatisation of PTTs and regulated competition, has increasingly played a role in developing remote surveillance and alarm collection systems.

As competition increases, and in-house resource decreases, the need to cut costs will be accelerated. More use

will be made of external contractors for maintenance and installation work, and more commercially available equipment will be purchased.

Energy constraints will require more efficient plant and more developments will be targeted at reducing the standing loads on power plant. This will be particularly applicable in new technology areas, such as telecommunications over passive optical networks (TPON).

This article describes some of the developments that have taken place in the past and suggests factors that will impinge on power and BES for telecommunications in the future.

History

When telecommunications was first introduced, the equipment was provided in basic accommodation. The earliest telephone systems were entirely manual, and little consideration was given to methods of providing AC or DC power supplies. As today, batteries were the prime power source, although in some locations local Electricity Boards provided a mains DC power supply. As the demand for more reliable and automated telecommunications equipment developed, so did the need for better power and cooling plant.

Twenty five years ago, telecommunications equipment was predominantly mechanical, with low heat

This article is based on a lecture given by the author to the North Downs Centre of the IBTE, and a paper presented at TELESCon 94 in Berlin

Figure 1—Example of central chilling plant

generation. Performance specifications existed, but they were not common throughout Europe and little was known of the effects of temperature and filtration. Today, the relationship between cooling, power and switch is linked by reliability, system requirements and energy consumption, but historically the developments have taken place independently.

Looking at these in more detail:

Environmental control

In the 1930s, heat dissipations of the order of 35 W/m^2 were common for mechanical switches and openable windows were considered adequate to provide cooling. Such systems had no methods of filtration, but where the dissipations were higher, mechanical ventilation with filters was provided.

Oil and dust particles were known to affect the maintenance of mechanical equipment, and ventilation systems became equipped with disposable or electrostatic filters. The disposable filters were predominantly in panel form, but in the large exchanges electrostatic filters were a common choice. In the mid-1970s, deterioration of the rubber mounts on uniselectors was attributed to ozone generated from the electrostatic filters. Investigations proved that ozone levels exterior to buildings were often higher than those inside, and that it was the type of rubber used on the mounts at fault. However, concern over safety issues associated with 10 kV power supplies, and complaints received from the public about the discharge noise associated with electrostatic filters, caused them to be removed from the network.

These investigations raised many issues about filtration and, after extensive work, new standards for equipment rooms were set.

In the late-1970s, reed-relay exchanges with higher heat dissipations were introduced, and required air movements up to a maximum of 16 air changes per hour (ACH). Semi-electronic exchanges further increased the heat gains, and



well-defined cooling policies were required. 16 ACH was adopted as a maximum figure to cool loads up to 270 W/m^2 and BT embarked fully on the road of using fresh-air cooling as a priority. Dual duct and fan systems were installed with the outlets from one supply interspersed with those from the other. Failure of one air delivery system therefore did not totally jeopardise the switch, since it was unlikely that both supplies would fail at the same time.

Risk calculations introduced new dimensions into power and cooling services. Evaluation of meteorological data showed that fresh-air cooling would suffice for 75% of the year and would give reliable and economic cooling, with refrigerated chilling only required during warm or humid conditions. By subdividing the load among a number of discrete cooling units, the site reliability was increased, and by adding redundancy, the reliability was further enhanced.

In the 1980s, a 15 kW unit was developed to control the apparatus room to $24^\circ\text{C} \pm 3^\circ\text{C}$, although the range was later expanded to 5, 8 and 30 kW sizes. The unit was designed to operate predominantly on fresh air. It contained a direct expansion refrigeration compressor, condenser and associated controls on a removable inner section and had no external

equipment. Installation only required one common hole in the wall and the outer casing could be pushed in position, connected to an electrical supply, and the inner section rolled in place. The refrigerant chosen was R22, since it combined good performance with low carcinogenic risk, and the system had an early design of microprocessor control. A sketch of the design is shown in Figure 2.

Initially, air was distributed by a simple diffuser installed in conventional ductwork. However, the introduction of fully electronic exchanges produced higher heat dissipations, and it became obvious that it would be necessary to adopt the practices of computer centre design and use either a false floor or a false ceiling to distribute the air. False-floor systems were installed in a few locations, but predominantly perforated or slotted false ceilings were used. These allowed high volumes of draught free air movement, and development work at trial sites showed that installations with the slots at right angle to the rack layout gave optimum cooling.

Today there are nearly 10 000 AHUs in the field representing the largest single use of AHUs in the UK. Central installations using reciprocating, centrifugal, screw or absorption chilling machines exist, but they

Figure 2—Standard cooling unit

are in large sites that already had existing duct installations. Some chilled ceiling systems were installed in the late-1980s, but initial costs proved to be high and the design did not support a predominantly fresh-air cooling policy.

Heating

The increase in cooling requirements was matched by a corresponding decline in heating needs. Twenty--thirty years ago, many of the smaller buildings were heated with either underfloor or self-contained electric storage systems working on special, long-duration off-peak tariffs. They were, however, expensive to run, and difficult to control or repair.

To reduce energy costs, a controller was introduced that predicted the optimum switch on and switch off times. Running costs were reduced, but the introduction of shorter duration electrical tariffs meant that it was not possible to store sufficient electrical energy during the charge period. The systems declined, and fan-assisted electric heaters were installed in intermittently attended buildings and low-pressure hot-water (LPHW) systems at more frequently occupied sites.

As the equipment loads increased, and more use was made of in-room cooling, electric coils in the AHUs, became the predominant choice. Now, with cooling a more necessary requirement, nearly all AHUs, are supplied without any form of heating.

AC power

AC mains supplies in the UK have always been considered reliable by comparison with many other countries, and in the infancy of telecommunications, the reliability of mains was not considered a problem.

AC prime movers were first installed in UK telecommunications in the 1920s for use in remote repeater and radio stations. The sets used single-cylinder compression-ignition engines capable of starting from cold, although 'cold' was not clearly defined. They ran on diesel

fuel, but could also run on paraffin or furnace oil in an emergency. They had especially large flywheels to overcome deficiencies in engine speed control, and typically drove two DC generators producing a continuous 40 kW. These first sets were therefore not designed to act as a stand-by engine, but to generate direct DC voltage.

Generally in the 1940s and 50s, exchanges were still not equipped with prime movers, as most of the sites were in town areas and could utilise the public electricity supply for battery charging via rectifiers. They therefore relied totally on battery reserve in an emergency. In the 1960s, two factors led to a reappraisal of AC generating plant, and its scale of provision. Firstly, a policy required that all communication installations should be capable of operating for at least 5 days in the event of failure of the public supply. Secondly, transistorised transmission equipment, with an inherent constant base load, was rapidly being introduced, and it became uneconomic to install batteries to supply a 5 days reserve. A decision was taken that sites with over 600 lines capacity should be provided with AC stand-by plant and this necessitated a reappraisal of the design considerations. The requirement of the sets were that they should:

- have high-reliability controls,
- be modular in design,

- be produced in a series of nominal sizes, and
- be designed to reduce on-site installation time.

All of the sets were water-cooled with automatic make up facilities, auto-start capability and a comprehensive visual display panel. The plant had four major modular elements consisting of the engine control cubicle, engine generator set, starter battery and emergency switch. The modular approach allowed rationalisation of the components to take place; for example, the control cubicles were made in 60, 100, 200 and 300 A/phase sizes to cover a wide range of plant. The modular approach also allowed easy installation and enabled different manufacturers to supply different items of equipment.

The control system monitored the mains supply and would start the generator if the voltage fell by more than 10%. Restoration of the supply voltage to within 7% of its nominal voltage for more than 52 seconds, would automatically switch back the mains supply. Initially, a 5 second delay was provided before an automatic start sequence was initiated, but this was later increased to 30 seconds to reduce cycling.

All of the sets were installed in separate plant rooms and environmental control was provided by automatically operated louvres. The sets were installed in 15, 31, 44, 72, 110, 145 and 216 kVA sizes (see





Figure 3—Power Plant 426

Figure 3). This provided a range of plant suitable for the medium size exchanges, but the growth of larger exchanges required a different, higher output design.

The physical size of the larger sets necessitated a modular approach to enable the plant to be installed in existing buildings and the sets were designed with separate engine control, contactor and switching cubicles, an engine generator set, and a separate compressed air or battery starting system. The larger sets were also the first to be fitted with bypass switches for the connection of a mobile generator set.

The first prototype was built in 1970, and the first sets commercially produced from 1971 onwards. Most operated with a mains supply of 415 V, but certain buildings had an 11 kV supply. For these sites, a standard packaged transformer was developed that further increased the system flexibility. The transformers were air cooled, rather than oil and negated the requirement for oil catchment pits and CO₂ extinguishing systems.

In the 1960s and 70s, a range of AC stand-by engines from 12–500 kW was available, but there increasingly became a need for stand-by plant for small exchanges. A series of sets covering the 4, 6 and 10 kVA sizes was therefore developed. These were simpler in design and used air-cooled engines, rather than the water-cooled



units used in larger installations. The air-cooled format also reduced costs, and, for the first time, engine speeds up to 3000 rev/min, rather than 750 or 1000 rev/min were used.

No major changes took place until 1984, when two important developments took place. The first of these was to combine the engine and its control gear in a single cubicle, as there was an increasing need to install the stand-by plant in the same accommodation as the switch, and negate the need for a separate plant room.

To achieve this, the engine was mounted directly above the alternator, with the motive coupling provided by a flexible toothed belt. This offered the additional advantage that engine speeds other than 1500 or 3000 rev/min could be used and opened up a wider choice of engine design. To reduce vibration, the engine and alternator were decoupled from the housing by a floating bedplate. The bedplate was resiliently supported on a secondary frame

which supported both the engine and the alternator. The second frame in turn was resiliently coupled to the cubicle. By this arrangement the primary mounts between the bedplate and secondary frame provided a high degree of isolation from the lower frequencies, while the secondary mounts between the frame and cubicle attenuated the higher frequencies. To reduce the acoustic sound further, a silencing system was specially designed to achieve a noise level of 75 dBA at 1 m.

The set was initially produced in a single cylinder 6 kVA rating, but to improve reliability, a twin-cylinder version was later introduced. The range was later increased to 10 kVA and 15 kVA sizes, all using the same housing.

The second development was an on-board microprocessor control system, referred to as *stand-by engine microprocessor automatic control* (SEMAC). This represented a major step forward from the earlier use of signalling relays, ratchet pulsing relays and relay logic. Each of these advancements in controls, represented, at the time, the latest state of the art technology, but SEMAC offered, for the first time, remote start/stop and interrogation facilities.

The system controlled the power plant via 32 analogue input channels, and a front membrane panel allowed the operator to take local manual control when required for test running and maintenance purposes.

Facilities available included:

- automatic 'mains fail' and local manual control;

Figure 4—2.8 MW 11 kV generating plant



- full monitoring and control of the plant;
- touch switch selection of automatic or manual control, including engine start/stop facilities;
- selection and display of power plant parameters;
- continual software checks;
- remote interrogation, control and test running of the power plant; and
- a history database of faults.

All of the facilities available on the front display panel were available to a remote user via an operations link, and it was possible to alter the system parameters locally or remotely via a PC or dumb terminal.

For the first time, an engine control system existed that could be incorporated into an overall control overlay, and the new 'COMPAC' design of set, with the new controller was introduced into the field in 1984. SEMAC was also fitted as standard to the complete range of fixed engine sets, with the exception of the large 500 kW and above sizes, that used a programmable logic controller.

In more recent years, most developments have been directed at the smaller size of set, one example being the introduction of a 'balanced flue' design in 1989.

This concept was introduced to reduce the footprint size of the engine set and used only one wall face and louvre for both inlet air and exhaust. By arranging the controls on a slide-

Figure 6—500 kW mobile gas turbine




Figure 5—Growth of stand-by plant

out cabinet frame, and the engine and associated equipment so that they could be maintained from one side only, it was possible to position the unit in the corner of the accommodation and save 30% floor space. So successful was this concept, that the range was later increased to cover 20, 31 and 44 kVA sizes.

There now exists a type and size of set that can cater for most installations, and currently there are approximately 5300 sets in service in BT. Figure 5 shows the population profile since 1964.


DC power

Considerable changes have taken place on the DC side since centralised

supplies were first introduced in the 1900s. Early plants used manually controlled motor generator sets to charge lead secondary cells, and two batteries were used enabling one to be charged while the other was discharging, thus giving rise to the term *charge discharge working*.

In the 1960s, DC motor generator sets were no longer installed, and the DC plants developed used mainly semiconductor rectifiers, coupled with lead-acid batteries. These were float charged continuously, and discharged only when the output float voltage fell beneath a predetermined level. The voltage range over which the majority of telecommunications equipment

Figure 7—Part of a colourful past—A metal arc rectifier



worked was from 44 to 52 V, and allowing for a 2 V drop in the distribution system, this set the power plant voltage limits from 46 to 52 V. The limits were derived from the earlier two battery charge/discharge systems and were based on a discharge voltage range of 2.06 V to 1.84 V/cell for a 25 cell lead acid battery.

The limits were not ideal for automatic working because of the small voltage difference between the maximum and minimum values, but it was not cost efficient to alter the voltage range, and apart from the later introduction of wide voltage limit equipment, these limits are still in use today in the UK.

In the 1960 and 1970s, two main categories of plant were in operation for small and large exchanges. In the small typical 600 line exchange, the plant used a statically controlled potential rectifier, to float continuously a 24 cell battery at 2.3 V per cell. To prevent the exchange voltage exceeding the maximum limit, only 22 cells were normally connected to the exchange, but all 24 cells were switched into use during a mains failure. With this system, it was possible to maintain the supply voltage above the lower limit for a 24 hour period.

For larger installations, two types of plant were available with the larger size capable of supplying loads between 1000 and 20 000 A. The plant was installed in conjunction with stand-by engines, and used two or more paralleled 25 cell batteries which together were capable of providing the required reserve period. Cells were floated at a lower voltage than those used in the smaller 24 cell installations (that is, 2.06 V) and were allowed to fall to 1.83 V/cell during a discharge. The lower float voltage, however, meant that the cells lost approximately 1% of their capacity each day, and refresher charging was required at fortnightly intervals.

The electromechanical telecommunications switches in use at the time

caused the exchange loads to vary with traffic from a low standing load during the night, to two or three peaks in the daytime. The highest peak was taken as the criteria and the battery capacity sized to supply the peak load for one hour, although greater battery reserves were provided for sites not fitted with engines.

Traditionally, singleton battery rooms were installed in basement areas. In large exchanges with installations of the order of 20 000 A and switch accommodation on six floors, it was necessary to use large distribution systems to overcome excessive volt drop.

A reduction of the busbar sizes was achieved by siting the plant on intermediate floors nearer the load centre, and further savings were realised by distributing the batteries around the building.

While most telecommunications plant normally operated at -51.5 V, certain equipment necessitated a multitude of different output voltages; Figure 8 shows the array of standards and power plants that were possible in a multi-function centre. If the proliferation of different voltages had continued, the outcome would have been high ongoing costs for plant and accommodation, coupled with a reduction in system reliability, since failure of any of the central supplies would have resulted in a system failure.

To overcome this, in 1970 all new designs of switch (with the exception of transmission) were supplied with a single central supply of -50 V. Subsidiary supply voltages were derived from regulators, inverters or converters forming part of telecommunications equipment; the number and rating of converters were chosen so that failure of any one would not result in unacceptable loss of service. Later, DC systems also used fully-automatic self-contained supply modules, working with a stand-by rectifier and a 25 cell battery with end-cell switching facilities. This approach enabled the capacity of

power systems to match the exchange growth, and the modular approach embraced the whole power chain from the AC stand-by to the DC power plant and distribution systems. The advantages of the modular system were:

- a more common design of DC plant,
- greater flexibility in planning extensions,
- simple design to improve system reliability,
- greater flexibility in the utilisation of accommodation, and
- cost savings by the deferment of capital expenditure.

The last factor was one of the most important attributes of the early

Figure 8—Possible supply standards in large centre




Figure 9—Power equipment racks

modular systems, although later improvements meant that this type of plant was not installed after 1976.

Throughout the early years of DC development, use was made of Planté secondary cells with pure-lead positive plates and pasted negatives. High-performance thin-plate cells were used in the 1960s and standard cells for the 15–15 000 Ah capacity installations. For up to 2200 Ah capacities, enclosed cells in transparent containers were utilised, but above this size, open cells in lead-lined wooden boxes were installed. Considerable debate took place over the years, as to whether pure-lead positive plates with higher costs and weight provided better whole-life service costs than alloyed plates, but experience showed that by proper charging and regulation, a minimum, economically viable life of 20 years could be achieved with the pure-lead design.

Probably the most important developments in DC plant design were associated with the introduction of electronic digital exchanges, and resulted in the use of distributed power systems, recombination cell technology and a power management system with remote interrogation and control facilities.

Fast-track switch planning became an important requirement, and it became clear that shorter planning and installation times were required for the associated power. The existing DC systems generally required local battery and plant rooms, with sufficient spare capacity for a 20 year growth. However, by placing the newly designed rectifiers and a new design of battery in cabinets associated with the racks, and providing a direct AC feed to the end suite, it was possible to replace the centralised power plant and associated busbar system, thus eliminating the biggest cause of major service outage. Equally important, power systems could grow at the same rate as the switch, dedicated DC power rooms were not required, and systems could be installed with much shorter timescales.



To achieve this, special batteries were manufactured that operated on the oxygen recombination principle, and by limiting the products of electrolysis it was possible to install them in switch rooms adjacent to the switch. The batteries were suitable for rack mounting, and had a high energy-to-density ratio. Special switched-mode power rectifiers were also developed, that could easily be changed by non-specialised people. Multiple 28 A units were chosen to improve system reliability, and the design took advantage of the wider operating voltage limits of the electronic switches.

The other major development, was the introduction of a microprocessor controller that comprised a power supply unit, a central processor and a rack and external interface. The controller could be accessed via a dumb terminal or local PC, or via an RS 232 interface and formed an integral part of the overall surveillance system.

The controller had the following features:

- measurement of system voltage, current, and alarms;
- an on-board battery monitoring system and auto battery discharge facility;
- reduced rectifier load during engine run periods;
- on-board diagnostics and history data; and

- remote monitoring and control functions.

The mid-1980s therefore saw a radical change in the design philosophy of DC power. *Rack-mounted power* was introduced from 1984 onwards, and different versions were manufactured to satisfy different applications. One example was the need to replace or offload centralised power plant serving large areas of existing digital equipment, and this led to a variant that used rack-mounted power components in conjunction with the existing busbar distribution (see Figure 9).

Since the introduction of these systems, various improvements in components have increased rectifier reliability and currently cost reduced versions are being considered. Naturally ventilated 28 A modules are still the first choice, but larger fan-assisted sizes are being evaluated.

Today over 11 000 switched-mode rectifiers are installed in the network. Rectifier reliability has met the design requirement, but the life expectancy of the batteries has been disappointing.

Energy

In addition to advancements in technology, the energy crisis in the 1970s also had an effect on the design and operation of equipment.

The major impact was in the environmental control area and most of the initial savings were achieved by correct re-commissioning of systems and controls.

Subsequently, optimum start/stop controls were applied to both wet and electric storage systems, and higher efficiency boilers and some heat reclaim systems introduced. This era also saw the introduction of a proliferation of notices, urging thermostats to be turned down and lights to be switched off.

To maximise savings, energy managers were appointed to highlight areas of maximum energy usage, and buildings were allocated an *energy usage index*. Those with the highest index rating, or highest energy bill, were then investigated on a priority basis.

One identified area of high energy expenditure was lighting, amounting to 18% of the energy bill. To reduce this, lower illumination levels were specified and low-energy fluorescent tubes installed, saving 10% per fitting. Equipment rack lighting was specifically targeted and various control schemes introduced, from simple timers, to systems using infra-red or ultrasonic sensors to detect people movement. These systems were introduced in the 1980s and generated considerable savings in lighting costs, since the lights only operated when work was carried out in the racks.

As energy costs rose, and the higher standing loads of electronic exchanges became more evident, the need to reduce the cooling energy bill further increased and fresh-air cooling became a more-important policy.

Attention also turned to more-esoteric schemes, including bore-hole cooling, where underground aquifers with water reservoirs at 12°C were used to provide free cooling. (See Figure 10.) By drilling to the aquifer, and installing two pipes a defined distance apart, water was pumped from the reservoir to cooling coils in room AHUs. Such systems were piloted at various locations, but the cost of boring the holes was prohibitive. Other systems, such as river-water cooling via heat exchangers, wind power, passive cooling and solar




Figure 10—Borehole cooling

panels, were also evaluated, but the UK climate, coupled with the capital cost of such schemes, made them mainly uneconomic.

On the DC front, rectifiers were purchased with improved operating efficiencies, and the number of spare rectifiers reduced to improve system efficiency. On the AC side, more efficient uninterruptible power supplies were selected, and engines installed that better matched the system load.

Future Trends

The design of power and BES equipment has evolved considerably since the introduction of the first telecommunications installations, and more advancements will take place in the future. The developments, particularly in the short term, will be driven by several factors including equipment costs, people resource, legislation, energy constraints and new technologies. Looking at these in more detail:

Costs

Electronics costs per customer connection have dramatically reduced in recent years, and as a result more attention is currently focussed on power and BES costs. Power costs, particularly at small sites, now feature as a significant percentage of the overall exchange installation, and considerable effort is being applied to reduce the cost of power.

One avenue is to consider more commercially available equipment. Care is necessary to ensure that the system requirements are clearly specified, but it is possible to achieve savings of the order of 50% by this approach on some equipment. Points to be considered are:

- Is a smaller alternator and engine acceptable, if there were a relaxation in the recovery time following a step load increase?
- Could, higher noise levels be tolerated?

- Are the sophisticated controls absolutely necessary, or is a five day, rather than a ten day fuel tank adequate?

Each factor needs to be assessed against the increased risks, but the use of proprietary items is likely to increase. Such an approach can also reduce the in-house testing requirements.

Opportunities similarly exist to rationalise the sizes and type of power and cooling equipment and thus gain economies of scale. Where one manufacturer supplies a range of equipment, value engineering can make better use of common components, although this may result in equipment being oversized for certain applications.

More can be made of longer-term contracts that can be re-assessed on a yearly basis and, by careful negotiation, manufacturers can be encouraged to take on more responsibility for providing technical back up to the field, rather than provide the service in house.

In Europe, the withdrawal of trade barriers can offer certain cost advantages, but prospective savings can easily be offset by increased travelling costs.

Reduced initial costs are not necessarily the optimum solution and increasing emphasis needs to be placed on whole-life costs. On-board self-diagnostics increase initial costs, but can significantly accrue savings over product life span. However, the necessity to recover costs in increasingly shorter times makes this approach difficult to justify, and only by in-house education will this dilemma be overcome.

Manpower resource

As competition increases, so will the pressure to reduce resource, and various statistics, including manpower per exchange connection, will be more carefully analysed. This will be true for the power and BES community; in some PTTs this is increasingly being considered as

non-core activity. Any reduction in resource will increase the need for remote functionality, and software will be designed to automate routining processes. All new major items of plant introduced into the network will need to be fitted with onboard controllers that will complement the control/surveillance centres. Routines, such as engine runs, that can be initiated manually from remote locations will increase, and systems will be introduced that can automatically routine plant without any manual intervention. Where changes in maintenance routines are introduced, software will become available that will download the information to the remote locations.

There will be an increase in work management systems containing databases of manpower resource and skill attributes. Systems will be introduced that will automatically contact the nearest person with the right skills, and more use will be made of small hand-held software tools to interrogate sites from remote locations.

A similar reduction in people resource in the field of switch and transmission will increase the need for multi-skilling. Switch maintenance people will be trained to carry out non-specialist work such as filter changing, visual inspection of engine sets, removing plug-in switch-mode power rectifiers. Specialist power and BES skills will only be necessary for functions such as changing batteries, routining flooded cells and fault finding on control cubicles. Coupled with this, there will be a significant increase in the use of contractors. This is already happening in the fields of lifts, refrigeration maintenance and portable appliance testing, but more contracts will also be placed for installation work. There will be an increasing tendency to include the cost of installation by the contractor in the contract price, and this will likely be extended to incorporate planning and site surveys as part of the package.

Turnkey contracts, where the contractor is asked to provide a complete ready-for-service installation of switch, AC, DC and cooling facilities, will increasingly find favour. Transferring design right responsibilities to manufacturers will similarly increase, although this needs to be balanced against any loss of intellectual property rights.

Legislation

Increasingly, as society's standards are raised so will new legislation be introduced.

Legislation can significantly impact on existing and future equipment and in the past decade several statutes and codes of practice have influenced plant design in the UK. These have included legislation on legionella, the COSHH regulations (Control of Substances Hazardous to Health) and the EAW (Electricity at Work) Act.

Currently documentation to reduce CFCs and HCFCs is increasing in importance and field trials are already underway to trial replacement refrigerants, but the existing replacement refrigerants will only be a temporary solution and further work will be necessary to develop alternatives that are environmentally acceptable.

Future legislation must be followed, but it is important not to overreact. The initial introduction of the Electricity at Work Regulations could have brought many DC and AC operations to a standstill, but careful analysis of the risks and discussions with the appropriate authorities made it possible to introduce acceptable working practises.

Such an approach should be adopted for all legislation.

Technological advances

Further improvements will be made in rectifier design efficiency, possibly to a point where fewer stages of rectification are required in the voltage chain. Energy costs will play an important part in the viability of TPON networks and concentrated



Figure 11—Power plant must work in all conditions

effort will be necessary to reduce the quiescent load of such systems.

The thermal efficiency of engines in stand-by plant will be increased, and the adoption of high-temperature ceramics will significantly improve engine performance.

More use will be made of fresh-air cooling and, if the design temperatures of equipment can be successfully increased, installations will be possible without any form of mechanical ventilation. Within the UK, some telecommunications equipment has already been successfully laboratory tested in ambients of 70°C operation, and many sites are already in service without refrigerated ventilation. Ultimately the goal should be to operate all telecommunications equipment without any form of refrigerated cooling, although significant advances will be necessary on the battery technology front to achieve this.

Conclusions

Considerable developments have taken place in the past fifty years on the power and cooling fronts, each one aimed at improving system reliability and performance. As competition increases, the necessity

to reduce costs and operate equipment with reduced resource will become more-important factors in the selection of equipment. Whatever happens, the adage **no power, no service, no revenue** will still figure high in the overall reliability of the telecommunications networks.

Bibliography

SPURGIN, D. A. Power Supplies in Telecommunications. BT Internal Publication.

Biography



John Parsons
BT Networks and
Systems

John Parsons is currently the manager of the Network Power, Cooling and Inventory Systems Unit, which has responsibility for all aspects of power and BES in BT Networks and Systems, including policy, operational support and budget build. He graduated from Loughborough University with in 1971 and has worked in various parts of the business since joining the company in 1972. Until recently, he was a member of the technical panels that write the guidance notes to the IEE Wiring Regulations and the Electricity at Work Regulations. He is a Chartered Engineer, a member of the IEE, a technical advisor to INTELEC and has presented papers at various international conferences.

TRIBUNE—A Broadband ISDN Test Bed

In a collaborative project as part of the RACE programme—Project TRIBUNE, a large test bed has been constructed to exercise the broadband integrated services digital network (B-ISDN) user network interface. This article describes the background to the project, and general aspects of the B-ISDN. The article goes on to describe the design of the test bed and its role in achieving the project goals. Finally a number of conclusions are drawn from the project.

The TRIBUNE Project

The TRIBUNE project is an element of the RACE Programme. RACE is an acronym for 'Research and technology development in Advanced Communications technologies in Europe'. The aim of this programme, part funded by the European Commission, is to set in place a common infrastructure for advanced communications services, in particular the broadband integrated services digital network (B-ISDN), within Europe. This first section of the article sets the scene for the TRIBUNE project, stating clearly the goals, the way in which the project is run, the scale of the work undertaken, and the background that gave rise to the project.

TRIBUNE project goals

TRIBUNE is an acronym for 'Testing, Ratification and Interoperability of the Broadband User Network interface'. Although somewhat lengthy at first reading, this expression does neatly encapsulate the essential project goals. There are two components to this project title. The first is the Broadband User Network interface (which is derived from B-UNI). This identifies this key interface as the focus of concern for the project. The second component of the project title states what is to be done to this interface; that is, to prove that this interface is correct and complete by ratifying its definition, by showing that it can interwork to other services and networks, and finally by showing how this interface can be thoroughly tested. In order to verify the user network interface in this way, the project has designed and constructed a test bed located in the research and development laboratories of the

Dutch PTT, near the Hague. The essential project goals are:

- to define, and implement, a specification of the broadband user network interface;
- to provide feedback to the standards bodies on the base recommendations;
- to construct a thoroughly tested test bed, available to third party users;
- to demonstrate the interworking between TRIBUNE and other network types; and
- to develop and utilise the necessary test methods and tools required by such an interface.

Project organisation

TRIBUNE is organised, like all RACE projects, as a collaborative piece of work. Representatives from each of the collaborating organisations form a consortium, charged with managing and running the project work, both individually and collectively. The TRIBUNE consortium currently has 12 partners, from seven European countries.

Project scale

The following parameters give an indication of the size of the TRIBUNE project:

Duration:	4 years, 1992–1995
Budget:	15.4 million ECUs (£10 million)
Manpower:	100 man-years
Test bed:	The TRIBUNE test bed occupies some 200 m ² of the Dutch PTT R&D laboratories.

Project background

The TRIBUNE project came about as the natural successor to an earlier RACE project known as the *BUNI Demonstrator*. The BUNI Demonstrator project was also concerned with the broadband user network interface (B-UNI), but at this time there was no agreed international standard(s) for this interface. TRIBUNE has built on the successes of the BUNI Demonstrator project by using the now available text from the ITU-T (formerly the CCITT) on the user network interface, by designing robustness and performance into the network side to provide a stable and reliable test bed, by the development of a class of purpose-built test tools, and by proving a wider range of interconnectivity types. TRIBUNE made good use of many lessons learnt from the BUNI Demonstrator project. In addition, a number of the BUNI Demonstrator consortium members continued into the TRIBUNE consortium; this continuity of personnel and key experiences in turn fostered a keen team spirit within such a large international project, which has contributed to the project's success.

Background on B-ISDN

This section discusses the rationale behind B-ISDN. It begins by comparing B-ISDN with the existing (narrowband) ISDN. It also discusses the so-called *asynchronous transfer mode* (ATM)—the mechanism defined to support the transfer of signalling and data between the user and the network. It concludes with a brief treatment of the B-ISDN standards.

Broadband ISDN versus (narrowband) ISDN

In order to understand some of what broadband ISDN can offer, we shall compare it with its predecessor, the N-ISDN. N-ISDN customers connect to their network operator via two 64 kbit/s data channels together with a 16 kbit/s signalling channel. N-ISDN is based on circuit switching

principles. When an N-ISDN user requires a connection, the network establishes a circuit and reserves it for that user. Since these circuits are at most 64 kbit/s in size, the services available over the N-ISDN are restricted (for example, telephony, fax, file transfer, etc).

B-ISDN users have an interface to the network that is 155 Mbit/s 'wide'†. The bandwidth of this interface has been standardised so as to accommodate many visual type services (TV, high-definition TV (HDTV), multimedia, etc). When a B-ISDN user makes a call, a virtual circuit (VC) is established; the user can specify the nature of this circuit; for example, its direction, bandwidth, quality-of-service characteristics, etc. A user may have many different connections established simultaneously, provided that their total bandwidth does not exceed the maximum allowed. Not only will the B-ISDN allow a much greater (and more flexible) range of services owing to the increased bandwidth, but it is inherently able to support more sophisticated connection types (for example, point-to-multipoint connections).

ATM—the transport for B-ISDN

Asynchronous transfer mode (ATM) has been chosen as the technology for delivering B-ISDN services. It is also built into the switching fabric of the network. ATM packages all data into fixed-length packets, referred to as *cells*. Widely different information streams can be multiplexed onto a single ATM interface, the separate streams being labelled by the virtual circuit identifier within the cell 'header'. ATM can support both synchronous (for example, telephony) as well as asynchronous (for example, fax) services.

† In fact, other such bandwidths are possible, for example 34 Mbit/s or 622 Mbit/s; TRIBUNE uses only the 155 Mbit/s interface.

Within a B-ISDN network the switching is carried out on individual cells. Contrast this with a circuit-switched network in which a complete (typically 64 kbit/s) circuit is switched for the duration of a call. In a B-ISDN switch, the cells associated with a given connection are only switched to the output as and when they arrive at the switch. This multiplexing of users' cell streams makes better use of the switch as a result of its statistical nature.

B-ISDN standards and the B-UNI

This section describes the status and progress of the international standards on B-ISDN, in particular as they relate to the B-UNI. Standardisation of the B-ISDN is under way within a number of recognised standardisation bodies (for example, the International Telecommunications Union (ITU), European Telecommunications Standards Institute (ETSI), ATM Forum, etc). These standards cover a whole range of topics covering general overviews, service descriptions, network architectures, user interfaces, network interfaces, to name a few. It is important to realise that these documents have not been invented completely from scratch; many of them can trace their ancestry back to the corresponding N-ISDN standards (indeed many B-ISDN standards committee members helped to develop the N-ISDN counterparts).

One of the central pillars of the B-ISDN concept is the interface between the broadband ISDN network and the user (that is, customer); this is the B-UNI. Its specification is a key element within the complete set of B-ISDN standards, and it forms the central point of focus of the TRIBUNE Project. Because of its importance in the B-ISDN world, we shall pause here to describe an important aspect of the B-UNI, known as the *protocol reference model* (PRM).

The PRM is a tool that is used to aid the description, specification, implementation and testing of the B-UNI. In order to discuss the complex-

Figure 1—Outline design of the TRIBUNE test bed

ity of the B-UNI, its behaviour is split up into a number of layers, each of differing complexity†. The lowest layers have the simplest of functions and provide a service to the layer above them. The highest layer provides the most sophisticated or complex functions (and is correspondingly harder to specify). The PRM describes the interaction between the network and the user in the form of protocols, or dialogues, which operate over the layers of the PRM. A fuller description of the PRM is given in Appendix 1.


The TRIBUNE B-UNI

The TRIBUNE B-UNI specification was derived primarily from ITU Recommendations, approved documents where possible and draft recommendations where no fully agreed standard was available. At the time when the original specification was being compiled during late 1992 and early 1993, the ITU Recommendations were incomplete in a number of areas, particularly the ATM adaptation layer and access signalling. In these cases, additional information was taken from relevant ETSI material and ATM Forum specifications.

Considerable effort has been made in compiling the TRIBUNE specification to make it complete, consistent and implementable. The majority of areas left 'for further study' in the standards documentation have been resolved.

A major objective in developing the TRIBUNE specification was that it should be *implementable*. This requires that the text be consistent, but also requires that it be understood as easily as possible. Information has been added to several parts of the document to illustrate the dynamic characteristics of its operation. A key aspect of this is the

† A similar approach is used in the N-ISDN standards, using three layers to describe the UNI.



inclusion of overview message flow diagrams which were included to explain clearly the operation of the protocol. These are easily interpreted manually, giving an unambiguous, readable description of the protocol.

During the implementation of the specification, it was found that the text taken from ITU Recommendations could in many cases be clarified to aid correct implementation. The recommendations are open to a degree of interpretation that is not suitable for a project such as TRIBUNE. A number of standards contributions were made, and modifications made to the TRIBUNE specification to specify clearly what was required from implementations.

The TRIBUNE Test Bed

This section describes the TRIBUNE test bed, and how it has been used within the project. The decision by the project consortium to design, construct and operate a test bed was taken early on. TRIBUNE is essentially a pragmatic project; having developed the specification of the B-UNI, it was considered essential to demonstrate that such a paper definition would work in practice. In the opposite direction, practical feedback into the B-UNI specification from the test bed was anticipated. A number of objectives were set for such a test bed, the most important of which are listed below.

- The test bed should be quick to design. At the start of the project the expected duration was three years. To progress from initial project goals to a complete integrated test bed with public demonstrations within such a

period, all within the context of an international, collaborative organisation, was a significant challenge.

- The test bed was to serve as a reference model. From the beginning, it was required to use the TRIBUNE test-bed to prove a range of interworking scenarios. In addition, it was hoped to make the test bed available to as many user projects as possible. By widening the audience for the TRIBUNE B-UNI specification, a greater degree of feedback from implementors would be achieved.
- This last aspect of the TRIBUNE test bed, its use as a third-party testing facility, is a double-edged sword. Together with the advantages gained by such an approach, there are also constraints placed on the test bed. The first of these concerns robustness. In order to support a wide selection of equipment types being connected to the TRIBUNE test bed, it must be assured that the network side can run unattended in a continuous and reliable fashion. It was also considered important that the test bed provided a reliable service to the 'users', and one within the expected performance margins agreed within the project.

The design of the complete TRIBUNE test bed (both the user and network sides) is now described. Figure 1 shows an abstraction of the TRIBUNE test-bed design, identifying merely its essential functional blocks. Figure 3, in Appendix 2, steps down a level of detail, showing the standard representation of the test-bed design.

The simplified TRIBUNE test-bed design shows the complete system as comprising four essential parts:

- *The switching 'fabric'* Within TRIBUNE there are two separate switches. One of these is a synchronous digital hierarchy (SDH) cross-connect, used for the distribution of TV and HDTV signals. The other is the ATM VCI/VPI switch (virtual channel identifier/virtual path identifier); this represents the core of the network switching and transmission function.
- *The test-bed control network* This collectively represents the central B-ISDN exchange intelligence. At the centre of the control network is the TRIBUNE control server (TCS). This element is responsible for the control of calls across the TRIBUNE test bed, coordinating the actions of the other controllers as required in the establishment, modification and release of all the calls through TRIBUNE.
- *Gateways* These extend the range of TRIBUNE by providing access to other network types. For example, one of the gateways has connections to the Dutch public N-ISDN; another one is connected to the European-wide ATM field-trial network which is currently undergoing experimental trials.
- *The users* This term refers to the complete collection of user terminals and customer premises equipment (CPE) that connect to TRIBUNE over the published B-UNI.

The above description gives a very brief description of the test-bed design. Readers interested in a more detailed technical description of the test bed can refer to Appendix 2.

Designing and constructing the test bed was the second stage of the TRIBUNE project (the first being the production of the B-UNI specification). The next three sections describe



how the test bed has been used to evaluate the B-UNI specification.

Testing the TRIBUNE B-UNI

TRIBUNE developed from an earlier (RACE I) project, known as the BUNI Demonstrator (see earlier). Two major problems found during integration of the BUNI Demonstrator were the time and effort required to debug firstly the user-to-network interface, and secondly that part of the test-bed control network concerned with access signalling and the call control functions.

The second of these was (relatively) easy to resolve in TRIBUNE. The interface between the network side access signalling and call control functions was redesigned to be optimised and simplified. This improved design did indeed pay off in greatly reducing the time and effort required to integrate these two control elements of the test bed (the new design also increased the performance of these two functions).

The first problem, that of integrating the network and user sides in a reasonable time-scale required a completely different approach to the design and implementation of the access signalling software in both the network and user sides. What was discovered during the testing of the BUNI Demonstrator project was that the (apparently) unambiguous specification of the access signalling layer in fact hid a number of inconsistencies. In many cases these did not become apparent until testing was performed. This difficulty with the specification had to be addressed in TRIBUNE.

A three-pronged attack at this problem was developed. The first weapon was the specification itself. Early on in the TRIBUNE project it was agreed that in addition to the traditional text specification of the two upper layers of the B-UNI, a more rigorous definition was also required. The method selected for this approach was to use a 'notation' referred to as specification and description language (SDL). SDL is a means of specifying in a more formal manner the actions of a system using a graphical representation of the system under consideration. Several tools are now available which will generate source code from SDLs. Furthermore, such tools allow the SDLs to be used to generate test cases required to exercise the protocol.

A second important innovation for TRIBUNE was the use of the so-called 'glass box' testing method. This technique includes the requirement for standard points of access to the software being tested by the test equipment. These are referred to as *points of control and observation* (PCOs), a term borrowed from the chip testing technique (as far as we are aware, TRIBUNE is the first large project to use the glass box test methods in a real application).

The third line of attack at the B-UNI testing problem was the development of a range of test equipment collectively referred to as *BITs*—broadband interface testers. A variety of BITs have been developed for testing specific layers of the network and user sides of the B-UNI protocol stack. The BITs connect up to the

relevant protocol layer through the PCO so that the various layers of the protocol stack can be tested individually and separately (an important capability, since the entire protocol stack was developed over several sites).


The use of formal specification methods for the B-UNI, the adoption of the glass box test methods and the comprehensive use of the family of BIT testers has helped to greatly speed up the integration of the user and network sides of the TRIBUNE UNI, and furthermore has greatly enhanced the manageability of the overall testing process.

Ratifying the TRIBUNE B-UNI

The process of ratifying the B-UNI is somewhat less well defined than simply testing it. Ratification goes beyond the mere static testing of the definition of an interface, by using real services over the interface. This service level testing will exercise the dynamic characteristics of the interface definition.

This ratification process formed an important part of the test-bed system-level integration, effectively the final stage before proclaiming the test bed as 'ready for service'. A public demonstration of the consortium's level of confidence in the project was shown in two open days at the end of September 1994. In excess of 60 engineers and managers came to see what TRIBUNE had produced, and an important part of this process was the demonstration of a wide range of service scenarios.

During 1995, the test bed will be enhanced in an important way to demonstrate further the suitability of the B-UNI for wide-ranging service support. The test-bed switch and control system will be upgraded to support the other 'flavour' of broadband UNI, the so-called ATM Forum UNI. The ATM Forum interface is the private version of the public UNI which TRIBUNE currently supports. Because of its commercial backing, in particular from many equipment suppliers, this private UNI is expected to develop a large market in a



relatively short interval of time. The operation of services between an ATM Forum and TRIBUNE B-UNI will be a further ratification of the TRIBUNE specification.

In addition, a variety of disparate users, many of whom have developed B-ISDN terminal equipment as part of separate projects (RACE and otherwise) will extend this ratification process by connecting to the test bed during late 1994/1995 for testing purposes.

Interworking the TRIBUNE B-UNI

TRIBUNE does not exist as a broadband ISDN test bed in splendid isolation from the rest of the world. It was recognised early on in the project that it was necessary to show that such a network node could interwork with other network types. Testing and demonstrating typical commercial services over these external connections formed an important part of the B-UNI ratification process, and in particular proved its ability to interwork with other disparate networks. TRIBUNE has three such types of interconnection, which are now described in turn.

- **N-ISDN** The TRIBUNE ISDN gateway provides access via two N-ISDN lines to the Dutch public N-ISDN. N-ISDN based services are implemented on two multi-service terminals (MSTs) connected to TRIBUNE via a special adaptor unit. The MSTs can therefore reach any terminal worldwide that can be accessed via

a combination of N-ISDNs/PSTNs. For example, during the open days, visitors were invited to make a telephone call from either of the MSTs to a colleague back in their home country

- **Local area networks** To ensure that TRIBUNE is able to interwork with this important network type, a proprietary LAN is connected to the test bed via an interworking unit. The project selected an FDDI LAN.
- **Broadband test beds** Throughout Europe, a number of similar broadband prototypes/test beds are being developed as part of the RACE programme of research and development. In order to connect to these other 'islands', TRIBUNE is equipped with two medium adapters. One of these operates at 2 Mbit/s, the other at 34 Mbit/s. The 34 Mbit/s medium adapter is connected into the European ATM pilot network and, through this, can gain access to the other connected broadband islands.

Summary

This concluding section highlights the key achievements of the project, and the lessons that have been learnt. The production of the TRIBUNE B-UNI specification represents a major project goal in itself. The paper version of the specification is a substantial document—approximately 600 pages. Its pragmatic form, together with its wide distribution

across Europe, are a feat in themselves. The design, construction, testing and operation of the TRIBUNE test bed to support this interface is the most visible achievement by the project team. This process has resulted in valuable practical feedback on the B-UNI, some of which has been in the form of contributions to the international standards bodies. The promotion of the test bed and its availability to third-party users has benefited both the users, and the TRIBUNE project itself; each such additional user represents a further step in the process of testing and ratifying the B-UNI definition.

An important lesson learnt from this process is the amount of time and effort that is required to develop a complete, unambiguous and implementable B-UNI specification. There is a considerable difference between an interface definition that is purely a paper exercise, and one which is enshrined in real hardware, software and human operators. The project aimed from the beginning to develop an appropriate set of test tools and methods, and this has paid off in an essential way during the test-bed integration process.

With an eye to the future, the TRIBUNE test-bed will incorporate the ATM Forum User Network Interface during 1995. In addition, operations and maintenance aspects of the interface will also be included, as well as experiments on a variety of important performance parameters. Beyond 1995, there are plans for the TRIBUNE test bed to be used in the next phase of the RACE programme, known as ACTS—advanced communication technologies and services.

Appendix 1—The Broadband ISDN Protocol Reference Model

As noted earlier, the B-ISDN PRM is divided into a number of layers which together provide the complete communication service between the user and the network. The protocol reference model employs a layered

approach to provide the complete set of communication services required over the B-UNI. A simplified representation of the B-ISDN PRM is shown in Figure 2. There now follows a description of the various parts of the PRM, what they do, and highlight some of the elements of the complete PRM not shown here.

The first point to note is the separation into two planes, the user plane and the control plane. The user plane, as its name implies, is responsible for the users' data (in whatever form it appears—voice, text, video, file data, ...). The control plane is responsible for handling the signalling between the user and the network, which is used to establish, modify and release calls. Both planes share the same three lower layers. The control plane has a fourth layer, the access signalling layer. These are now described briefly.

Physical layer

This lowest layer of the PRM is responsible for the physical transmission of the individual bits that comprise the user or control data sent over the user-network interface. The physical layer can cover a number of options; for example, electrical versus optical lines, SDH framed ATM versus 'pure' ATM. In addition, the physical layer inserts and extracts idle cells to maintain cell delineation.

ATM layer

It is in this layer where the knowledge of ATM cell formats is held;

every ATM cell is 53 octets in length, comprising a 5 octet header (containing routing information), together with 48 octets of users' data. The sending ATM layer will ensure that correct headers are attached to the users' data, and will also multiplex the cells corresponding to the various virtual circuits that can coexist over a single user-network interface. The receiving ATM layer will receive the cells passed from the physical layer, and demultiplex the various virtual circuits.

ATM adaptation layer (AAL)

The AAL is used between the data stream and the ATM layer on the sending side by segmenting data into ATM cells; at the receiving AAL, it reassembles the incoming cells into a continuous stream of data for the layer above it. The AAL also performs some error detection and correction functions. The AAL comes in a range of types. AAL1 is optimised for the adaptation of synchronous data; for example, speech or video. AAL types 3/4 and 5 provide adaptation for asynchronous services; for example, file transfer. AAL type 5 is a faster, and more 'lightweight' protocol than 3/4. AAL type 5 has a variant that is selected especially for the transfer of signalling messages.

Access signalling

At the top of this protocol stack in the control plane sits the access signalling layer. Within this layer are defined all the procedures for accessing, monitoring and controlling the wide range of B-ISDN services. The user (and the network) will use access signalling messages to establish calls, to monitor the state of a call, to change the characteristics of a call, and of course to release a call. Access signalling also contains recovery procedures for restarting a connection, or the entire user-network interface.

This description has shown the essential ingredients of the B-ISDN PRM as required for this article. There are two other essential

Figure 2—B-ISDN PRM




Figure 3—Overall design of TRIBUNE test bed

ingredients of the model which are not shown here. These are the plane-management function, and the layer-management function. Plane management coordinates between layers, and between the user and control planes, while layer management is responsible for the management functions within a single layer.

Appendix 2—Overview of the TRIBUNE Test bed

This Appendix includes a more detailed description of the various elements of the test bed. The 'standard' representation of the test bed is shown in Figure 3.

Switching fabric

The test bed comprises two distinct switches. The SDH switch is in fact a simple cross-connect, whose purpose is to provide switching (and distribution) of TV and HDTV signals. The main test-bed switch is an ATM switch. This is able to switch cells on a given input VC to any nominated output VC value. Each of the switches has its own dedicated controller, a PC in the case of the SDH switch, a workstation for the ATM switch.

Test-bed control network

This is often referred to as the central nervous system of the TRIBUNE exchange. It houses all the control and coordination functions that are needed to run the network side of the test bed as a whole. The TRIBUNE control network (TCN) is made up of the following elements:


- **Two switch controllers** These devices are supplied as part of their associated switch, but form a part of the TCN.
- **Network signalling termination (NST)** This device is the point in the network in which the B-UNI access signalling protocols are processed. The ATM switch is pre-configured so that all signalling cells from the users are automatically routed to the NST.

- **TRIBUNE control server (TCS)** The TCS sits at the centre of the test-bed control network. The single most important function of the TCS is the call control function. On receipt of signalling messages from the NST, the TCS will coordinate the actions of the rest of the TCN in order to establish, modify and release calls across the TRIBUNE test bed. The TCS interacts with the other elements of the TCN by the use of a specially designed set of messages transferred over a LAN, the TRIBUNE control LAN.
- **TCMS testbox** This is a PC-based tester designed to be able to test each of the separate elements of the TCN in isolation (that is, before integration has commenced).

Test-bed gateways

There are essentially three such devices:

- **N-ISDN gateway** This provides access to the N-ISDN, and is itself an element of the TRIBUNE control network. The TCS exchanges messages with the N-ISDN gateway to process calls arriving from, or destined for, the narrowband ISDN. The N-ISDN user data (that is, the B-channels) is passed to the ATM switch via an adapter unit known as the *system access interface (SAI)*.
- **LAN gateway** Referred to as the *LAN interworking unit (IWU)*, its job is to translate between the protocols used over the FDDI LAN and those defined over the TRIBUNE B-UNI by the protocol reference model.



- *Broadband gateways* There are two of these; their only essential difference being in data rates, one operating at 2 Mbit/s, and the other at 34 Mbit/s. As their name implies, they act at the physical layer of the protocol, and provide access to the ATM field-trial network and the so-called *Megaswitch network*.

Users

One of the most important ingredients of the project! As wide a variety of user types as possible have been encouraged to connect to and use the TRIBUNE test bed. The users shown below include:

- *TV and HDTV equipment* This includes a variety of TV and HDTV sources plus monitors. In addition, the HDTV signals are passed through Eureka 256 coders/decoders.
- *Narrowband ISDN multi-service terminals (MSTs)* These are modified commercial products that offer the usual range of narrowband ISDN services. Between the MST and the ATM switch sits an adapter unit which maps between the TRIBUNE B-UNI and the N-ISDN interface presented by the MST. Using the two MSTs it is possible, via the N-ISDN gateway, to operate N-ISDN services across the test-bed to any other terminal accessible from the Dutch public N-ISDN.
- *Windows PCs* The use of 'standard' PCs as a platform for B-ISDN terminals is expected to play an important role in the widespread usage of B-ISDN services. The PCs can run a range of applications, and communicate with the TRIBUNE switch via an ATM interface card.
- *Broadband terminal* This is based on a workstation and has been developed in line with a multimedia application, designed to exercise the flexible nature of

the ATM interface. The application involves the simultaneous transfer of moving pictures, sound, text, and still images over the B-UNI.

- *Broadband database* The database was developed in parallel with the broadband terminal and holds the associated multimedia data. This equipment element is based on a Sun workstation. The workstations in both the broadband terminal and database communicate with the TRIBUNE ATM switch via an adapter. This adapter unit is used in a number of configurations within the test bed.
- *Compliance testers* Two forms of the broadband interface tester (BIT) already discussed in the main part of the article are shown here. The ATM compliance BIT is designed to test the network implementation of the B-UNI ATM layer. The terminal BIT is designed to test the call control functions of the test bed. It is used by the operator in much the same way as an 'ordinary' B-ISDN terminal, but has a number of additional test-specific functions.

Acknowledgements

I would like to thank my colleagues Peter Hovell and Malcolm Marchant for their assistance in reading through the draft of this document.

Biography



Chris Shephard
BT Networks and
Systems

Chris Shephard began work for BT in 1980, initially on the development of software for the System X remote concentrator unit. He joined a small

team implementing a prototype version of the 1984 ITU Red Book ISDN standards. In 1986, He left BT to work for Siemens in the US. Here he was involved with the system level requirements and development of the Siemens public switching system (EWSD), adapting the Deutsche Bundespost Telekom product to the US Federal and local requirements. In 1988, he returned to BT to work on the control software of the BT-LA 30 Multiplexor. In 1991, he joined the BUNI Demonstrator team, and subsequently the TRIBUNE team. He is the TRIBUNE project manager, and leads several work packages within the project.

Glossary

AAL	ATM adaptation layer
A/D	Analogue-to-digital (converter)
ATM	Asynchronous transfer mode
B-ISDN	Broadband ISDN
B-UNI	Broadband UNI
BIT	Broadband interface tester
CPE	Customer premises equipment
D/A	Digital-to-analogue (converter)
ETSI	European Technical Standards Institute
N-ISDN	Narrowband integrated services digital network
ITU	International Telecommunications Union
LAN	Local area network
MST	Multi-service terminal
PCO	Point of control and observation
PRM	Protocol reference model
PSTN	Public switched telephone network
RACE	Research and technology development in Advanced Communications technologies in Europe
SDH	Synchronous digital hierarchy
SDL	Specification and description language
UNI	User network interface
VC	(ATM) Virtual circuit

BT Welcomes Telephone Preference Service

BT welcomed the launch of the Telephone Preference Service (TPS) as giving customers yet more control over how they use the telephone. The TPS is a service set up by the telemarketing and telecommunications industries to enable consumers to have their telephone numbers removed from lists used by the telemarketing industry. It is free to consumers and run by a non-profit organisation, the Telephone Preference Service Ltd.

Mike James, BT's Manager of Assistance Services, said: 'The TPS is absolutely in line with BT's commitment to responsible telephone usage and to ensuring that choice and control rest with the customer.'

'In recent years, BT has extended customers' choice and control with a whole range of services, from specialised malicious calls bureaux to caller display technology.'

'TPS will now help everyone use the telephone more efficiently and telephone sales people will be able to make more effective use of their time by calling customers who are interested in receiving such calls.'

BT customers who want to register with the TPS can call 0800 398893 for more information and registration forms or they can pick up a form from any BT shop.

Information Overload?

A computer program that offers effective answers to 'information overload'—where people have too much documentation to assimilate—has been developed by BT Laboratories.

The software includes an Information Agent, which searches texts stored on a computer to select those of most interest to the operator, and an automatic Summariser that can abridge text to a fraction of its length without losing significant detail.

BT has developed these techniques in response to the increasing mountains of information available worldwide and the ease of access to

it brought by modern telecommunications and computer networks.

The Summariser uses a statistical algorithm either to highlight the most important sentences in a text, or to discard unhighlighted sections and leave a précis. Any length of précis or summary can be chosen, from one to 99 per cent of the original article.

The basic technique is robust, can work on a wide range of text types and is independent of subject: equally good results have been achieved on topics as diverse as red squirrel behaviour and semiconductor laser physics.

Tests comparing summaries with authors' own abstracts have shown that an abridgement down to five per cent of the original length typically contains some 70 per cent of the original's important information—while a quarter-length summary keeps virtually all the information. BT believes millions could benefit from a reduction in paperwork of this scale.

Keith Preston, of BT's Natural Language Group, which developed the software said: 'There are many professions and business functions where fast, easy access to banks of data has brought huge advances in terms of availability, but there is a parallel disadvantage as well.'

'The fact that there is now so much 'on tap' means it threatens to overwhelm the very people who need to absorb it—whether they are doctors or scientists needing to keep abreast of the latest research, lawyers collating case material, or businessmen sifting through company and market information.'

Natural language processing

Natural language processing is a technology with the direct benefit of helping people to find and use information without having to sort the relevant from the irrelevant.

BT's Information Agent program manages databases of information and, over the time they are accessed by the user, builds up a profile of those subjects of most interest to that user. As items are added to the database the program automatically

arranges the contents into a priority order it has 'learnt' is of most interest to that user.

The key advantage over standard document retrieval techniques is its ease of use, which does not rely upon keywords or search strings. Instead, the Information Agent works on a 'find me more like those' instruction. The research programme underway at BT Laboratories is set to expand from dealing with pure text to cope with the challenge of multimedia, involving deeper analysis and comprehension of texts. It is moving toward building the intelligent network of the future which will not simply transport information but increasingly interpret, filter and process the information vital to specific professions and businesses.

BT Provides Business TV Links for Philips Event

When Dutch multinational, Philips Electronics, staged what is believed to be the largest ever, single, private business communications event in the world in January 1995, its IT services subsidiary, Philips Communications and Processing Services (C&P), turned to BT's Global Satellite Services (GSS) unit to provide the international satellite communications links.

GSS, which is part of BT's Global Communications division, provided satellite links for a one-hour long programme which was broadcast in 18 languages to over 200 000 staff at 600 sites in three continents—Asia, the Americas and Europe—on 9 January. The communications network involved the use of no less than four satellites, 2000 televisions, cinema screens, video-walls and superprojectors.

The satellite link-up launched Philips' 'Customer Day 1995' quality-of-service initiative, which included a keynote address by Philips' president Jan Timmer and a CD-i presentation on customer best practice. BT handled the broadcast signals to 38 cities, as far a field as Bombay in India, Los Angeles in the USA, Quebec in Canada and Oslo in Norway, with overall control for

Europe and Asia coming from BT's Business Television Service Management Centre in London and from Massachusetts for North and South America.

Low-Cost ISDN Package From BT

A new low-cost internetworking package, which combines the benefits of integrated services digital network (ISDN) and local area networks (LANs), ideal for small and medium-sized businesses, has been unveiled by BT.

The package, LANstar2, is designed to meet the growing demand for affordable, easy-to-use systems that support file transfer, LAN interconnection and remote LAN access—the fastest growing ISDN2 applications.

LANstar2, using ISDN 2 lines, allows the interconnection of LANs at office branches as well as providing teleworkers and remote LAN users with access to central host and database applications.

The package consists of a high-performance ISDN2 Ethernet bridge and KNX Access Server hardware and software. The KNX Access Server is a PC card which fits inside a standard 486 PC.

In addition to linking remote LANs, a company can connect up to ten standalone PCs, each equipped with one LANstar PC card, to an Ethernet LAN from anywhere in the world where there is access to ISDN.

The card also enables the allocation of up to four ISDN2 lines, giving increased bandwidth when required, such as moving large data files or video files.

Protection against unauthorised access to company information is provided by a range of security checks on the incoming call. This is known as *caller line identity password* (CLIP), and is even available through the MAC address of the calling ISDN2 card. In addition, LANstar2 offers a comprehensive filtering process that 'masks' local network nodes, preventing information from these devices being transmitted over the ISDN.

All Encompassing ISDN System

The need for numerous pieces of ISDN equipment and separate ISDN lines for voice and data communications has been reduced with BT's introduction of the Jtec Virtual eXchange.

Companies can now integrate networking equipment such as hubs, multiplexers, routers and voice networks (for instance PBXs and key systems) using ISDN2 and/or ISDN30 services. The Virtual eXchange can run multiple applications from one platform as it is an ISDN switching exchange in its own right.

Managing both ISDN and non-ISDN traffic, the system also optimises the use of a company's corporate network, whether it is private or public. As a result, various applications can run over one network, which can significantly reduce the cost of maintaining several circuits for different uses.

The Virtual eXchange supports numerous applications including, switching voice and data calls, routing of calls and associated numbers, ISDN back-up, bandwidth management and compression and ISDN to the desktop.

In the case of ISDN2 to the desktop, the Virtual eXchange can divide a single ISDN30 line into 15 ISDN2 lines once inside a building, eliminating the need to run multiple lines to a site. This is cheaper than installing numerous ISDN2 lines, a company's ISDN infrastructure is simplified and less ISDN2 terminating equipment is required.

The Virtual eXchange also aggregates up to eight ISDN2 or ISDN30 channels giving a maximum of 504 kbit/s of bandwidth, when required. This is ideal for applications like LAN interconnection or high-definition videoconferencing, which require more than 64 kbit/s offered by a single ISDN B channel.

The Virtual eXchange's modular architecture ensures a company has a fully future-proofed and flexible system that can support a single application now and be expanded at

any time to support multiple applications meeting business demands.

Bank Branch Network of the Future

The Royal Bank of Scotland has signed a contract with BT to upgrade its branch network communications infrastructure. The contract will involve changing the bank's entire branch network from analogue communications to a new digital link-up based on BT's most advanced systems.

On completion, every branch in Britain's most extensive branch coverage, from Shetland to Cornwall and from the Hebrides to East Anglia, will have a dedicated connection from its local area network (LAN) to the bank's main-frame computer.

This will allow the Royal Bank to deliver to its customers the most advanced banking information services available.

Norman McLuskie, managing director of the Royal Bank's operations division, said: 'This is the first step towards integration of voice, vision and data services for the benefit of our customers no matter where they may be. It is natural that BT, which has been closely involved with us on so many elements of our drive to create a nationwide British bank, should be our partner in this project.'

BT Sells Stake in Belize Telecommunications Ltd.

BT has announced that it sold its 25 per cent stake in Belize Telecommunications Limited (BTL), with MCI taking 23.5 per cent and the Belize government taking 1.5 per cent. The deal follows BT's alliance with MCI in which BT bought 20 per cent of MCI and established Concert—the BT and MCI global network company. Under the deal, MCI took responsibility for distribution of services in the Americas and BT is responsible for distribution elsewhere.

BT has held a stake in Belize Telecommunications Limited since 1988. During this period, the

network has doubled in size to around 30 000 lines, in a country whose population is around 200 000. BTL offers a wide range of national and international services and a cellular service was started in 1993.

Learn Your Euro-Manners

More than 20 years after joining the Common Market, Britain's business-people still believe they can forge across Europe with a blind disregard for foreign languages, conventions and etiquette.

A survey of 250 directors of major UK companies for Cellnet has found that, of those who travel regularly on business in Europe, around half admit they make no effort to study correct forms of local business etiquette and over three in ten say that they have not bothered to even learn a few phrases of the relevant language.

While one third of company directors admit that this lack of local knowledge had left them at a disadvantage in a business meeting with foreign counterparts, just one in five of British companies provide guidance for their travelling employees on how to behave and business etiquette abroad.

However, Cellnet has launched *Euro Manners—Cellnet's Guide to European Business Etiquette*. It is a light-hearted guide with a serious message, containing advice on everything from basic forms of address, meetings procedures, dress codes and body language, as well as specific advice on local business customs.

Euro Manners does not claim to give the answer to every question

about doing business in each European country, but one read should be enough to make people realise that they cannot simply jump on the Euro-train, arrive in Paris and expect to come home with their order-book full.

As a communications company Cellnet advises anyone using its mobile telephone service in Europe to at least learn how to answer their phone in a few languages.

Euro Manners is available free by calling 0800 424323.

BT and Apple to Jointly Market Mobile Data

BT and Apple Computer UK, have agreed to market jointly a complete mobile data solution designed to take the lead in the emerging mobile data market.

This alliance between BT and Apple will enable users of Apple's Newton MessagePad 120 to read their e-mail and access on-line information services using Cellnet's GSM cellular network. Users of Apple's PowerBook 500 series of portable PCs will be able to send and receive facsimile messages in addition to the other services.

According to Richard Gibson, data solutions manager, BT Mobile, 'Both the PC and mobile telephones have witnessed phenomenal growth, however businesses have tended to use them as separate productivity tools. This solution will provide mobile users with the opportunity to combine and exploit the very latest in mobile communications technology with leading mobile computer technology, at an affordable price'.

Tel-Me Service Agreement

BT and PhoneLink plc have signed an agreement to market and distribute Tel-Me, PhoneLink's new on-line business information service. Phone-Link has developed a customised version of Tel-Me for BT.

Tel-Me enables users to reduce radically the time spent on everyday business tasks. Currently Tel-Me users can, for example, find out train times from British Rail, quickly get journey plans with updated AA Roadwatch information, find and print vital company financial reports from CNN and Infocheck, and get up-to-the-minute news and weather information. Such vital everyday business information is available instantly in a common user-friendly format to all Tel-Me users.

Tel-Me will be marketed by BT as an ISDN solution, with BT's ISDN voice and data products. ISDN is of particular value for Windows applications which require the transmission of text and images at high speed and for users who need to transfer large quantities of data, text, voice and images easily and cost effectively.

BT will also offer a solution to work over the standard public switched telephone network, packaging Tel-Me with BT's new Prologue 1414 modem, which will include the Tel-Me fast-connect protocol. This allows a connection to be made much faster than using conventional modem technology, enabling a complete enquiry to be sent and the answer received within 5–8 seconds.

industry news

Interactive Digital Services Trials in USA

Southern New England Telephone (SNET) is scheduled to run a trial of a wide array of interactive digital service, including video-on-demand, enhanced pay-per-view, and 'cable like' programming in late 1995. The trial is planned to provide service to

a total of 150 000 homes in two Connecticut counties. The regional head-ends in each county will receive analogue satellite transmissions and centralised video server and applications services. A series of remote hubs will link application and video servers to end-users.

SNET will combine Sybase Intermedia products and SYBASE

System 10 relational database products for complete management of their interactive service business.

BellSouth has announced a forthcoming interactive television trial in Chamblee, Georgia. The trial will seek to provide a wide array of analogue and digital interactive services, including video-on-demand and interactive education. These will

be carried on a network that will pass 12 000 homes.

FCC Awards License for IRIDIUM Satellite System

At the end of January 1995, the Federal Communications Commission in the USA awarded Motorola Satellite Communications Inc., a license to construct, launch and operate the IRIDIUM satellite system. The satellite network will provide worldwide, hand-held wireless telephony services beginning in 1998.

The granting of this license is a significant regulatory action relating to global low-earth-orbit (LEO) satellite systems. The system will be the first operational wireless communications network designed to serve hand-held pocket-sized telephones with a single number virtually anywhere in the world. The telephones will be dual-mode, permitting users to interconnect with their existing cellular networks in their home market and the constellation of 66 low-orbit IRIDIUM satellites when travelling outside their home region. IRIDIUM services will include voice, data, facsimile, and paging.

Cut the Waste

An energy management initiative to encourage engineers to give priority to cutting the United Kingdom's energy costs is being launched by the Engineering Council, in partnership with the Energy Technology Support Unit (ETSU), for the Energy Efficiency Office.

A major challenge facing engineers is that of environment management, of which energy management forms a major part. It is estimated by the Energy Efficiency Office that 20 per cent of the nation's energy cost is wasted. Successful energy management can lead to significant cost savings as well as environmental benefits through reduced emissions.

The Engineering Council's code of professional practice on environmental issues highlights the

importance of energy management for engineers and recommends action to be taken.

The Council's project with the ETSU is to establish a continuing professional development training package for energy management. This will assist engineers to analyse their professional development needs in this area, and provide information on available learning resources and guidance on access to NVQ assessment.

The project is sponsored by the Energy Efficiency Office and will build on the national standards for managing energy, being developed and tested jointly by the Management Charter Initiative and the Energy Efficiency Office. The project is enthusiastically supported by the Institute of Energy and it is managed by The Engineering Council.

Material is planned to be available in late summer 1995.

Austrian Airlines Contract for Skyphone

Skyphone, the international passenger communications and multimedia consortium of BT, Norwegian Telecom and Singapore Telecom, has announced a half-million pound contract for passenger telephony with Austrian Airlines.

From April 1995, passengers flying to Beijing and Tokyo on Austrian Airlines' two new Airbus A340s will be able to make international direct-dialled calls to over 200 countries worldwide with the swipe of a credit card. Each aircraft will initially feature six bulkhead-mounted handsets across all classes. This will later be expanded to seat-arm telephones throughout business class.

PhONEday Arrives

On Sunday 16 April, all UK area dialling codes changed from '0' to '01' on all UK operators' networks.

Several other changes took place at the same time to increase the quantity of telephone numbers available. Five towns and cities received brand new area codes and

an extra digit in front of each telephone number. They are:

- Leeds (0113) 2xx xxxx;
- Sheffield (0114) 2xx xxxx;
- Nottingham (0115) 9xx xxxx;
- Leicester (0116) 2xx xxxx; and
- Bristol (0117) 9xx xxxx.

Also Kingston-upon-Hull added a pre-fix '3' in front of all five digit numbers. So a typical number such as 0482 xxxxx is now (01482) 3xxxxx.

The Learning Superhighway

A new collaboration between the private and public sectors will seek to create a learning superhighway by establishing multimedia learning centres in schools across Merseyside. February 1995 saw the foundation of MESH—the Merseyside Education Superhighway Limited. Its objectives are to promote cost-effective learning in Merseyside, across all ages, by utilising the latest technologies, content and support tools to deliver relevant education and training, and other services, locally. Multimedia learning centres in each school will be the hosts on a network connecting these locations together.

MESH resulted from an initiative led by the Liverpool City of Learning—using new technologies in education. The Liverpool City of Learning, launched on 14 December 1994, focuses on training and learning as keys to future prosperity in the region. It aims to make Liverpool a leading centre of excellence in learning practice and methodology.

During school hours, the centres will be used to support teachers in delivering the national curriculum and other qualifications. During these times, LINK staff will be on hand to support the teacher and support the network, technically.

Outside school hours, at weekends and during holidays, the centre will be used for commercial education and training. It is envisaged that each centre will provide education and training services to individuals,

companies and other organisations in its locality. These services will include national vocational qualifications, IT training and a range of other courses and skills.

ATM Link Breaks Record

In February 1995, a record was set for the world's longest asynchronous transfer mode (ATM) link, at the recent G7 Ministerial Conference on the information society. Extending from Burnaby, B.C., in Canada to Brussels in Belgium, the link was 11 000 kilometres.

The link, used to bring a distance education demonstration to Brussels, represented an interconnection of information highways and telecommunications links. The cross-Canada connection was made via the CANARIE national data highway—one of North America's longest single span links in the electronic highway infrastructure—to the EuroATM test network via Teglobe Canada's facilities. In Europe, it was routed through London to Brussels, the site of a demonstration showcase held in conjunction with the Ministerial Conference.

Software Helps Cut Cooling Costs

Flomerics Ltd. have been nominated by BT for the annual BT Network Product Quality Award for their


Flovent and Flotherm software programs. The award was first launched six years ago as part of BT's total quality management program, to encourage BT's suppliers to provide quality performance products.

The software programs are used by BT to analyse the movement of cooling air within and around switching and transmission installations, where large amounts of heat are generated by the electronics. The primary aim of the analysis is to reduce costs for cooling switching installations, by meeting the European Telecomms Standard, at minimum cost.

Flovent analysis of a sample switching installation predicted specific features of airflow which were later verified by actual measurements on site. For example, adjacent to the cabinets, in plan view, artificial test smoke was seen to separate and travel in two directions. The Flovent simulation predicted, not only the separations, but also their occurrence at different points in different gangways.

After further modelling, a simple cooling design was evolved which will have a significant effect on cooling costs. A whole-life saving of over 50 per cent is predicted, the capital cost of each cooling system has been halved and the skill level required for cooling system maintenance has been reduced, together with the man-hours required to look after it.

Temperature contours as predicted by Flovent software for a vertical section through a typical switching installation.



Geographic Telephone Numbering Scheme?

OFTEL Director General, Don Cruickshank, announced in March that he would shortly be launching a further consultation exercise on the new geographic telephone codes.

After PhONEday—all existing geographic telephone numbers begin '01'. This means that '02' to '09' will soon be available for future expansion. Industry consultation has indicated a preference for using '02' for a new geographic scheme.

At present OFTEL have a broad policy—no more than that—for the new numbers. The scheme might look like this:

- 01—current geographic public network; some paging services.
- 02—new geographic scheme.
- 03—unallocated.
- 04—mobile services: cellular, PCN; radio-based paging.
- 05—unallocated
- 06—unallocated.
- 07—personal numbering.
- 08—specially tariffed services: premium rate, national calls charged at local rate, freephone.
- 09—unallocated.
- 00—international access.

All the current publicity about how the numbers starting '02' will look is premature. The new ranges will not be introduced until they are needed and the details have been fully discussed with operators and customers.

Don Cruickshank said, 'Customers can be confident that, after PhONEday, no new changes to their numbers will be imposed on them. In the longer term, customers will have a choice between '01' and '02' for new numbers.

'However, there has been some concern about our (OFTEL's) plans for the '02' range. So that customers can have a full say in the future of their telephone numbers, we are going to put out our '02' proposals for consultation. We hope that everybody, and particularly customers, will let us know their views. It is very important that there are

enough numbers in the future for all the new operators and new services coming along. But we must get it right and customers must feel comfortable with the new scheme. We are aiming to publish a consultative document in May and complete the consultation exercise in September.'

Another Move Towards Number Portability

Don Cruickshank, Director General of Telecommunications is taking further action to promote the early introduction of number portability. This allows a customer to keep his or her number when changing telecommunications operators. He has issued, for public consultation, a modification to BT's licence which will enable OFTEL to determine the charges BT makes to another operator for number portability.

Convinced that the early introduction of number portability was vital for UK telecommunications customers, Don Cruickshank said that its absence is a major disincentive for customers to switch between operators and that its introduction would lead to benefits which all customers will enjoy.

The move follows disagreement between BT and other operators on the most appropriate way of dividing between them the costs of introducing number portability. Last August, OFTEL issued directions to BT to provide number portability to Videotron in the latter's cable franchise areas. However, although BT and Videotron were able to agree on the technical and practical aspects, they failed to agree on commercial terms.

Voice Processing Takes a Foothold

Results from a voice processing survey by Aspen Consultancy has shown that voice processing applications have already taken a strong foothold in the Call Centre environment.

The survey, *The Call Centre in Business*, revealed that more than 60

per cent of those surveyed had a platform for voice processing installed, either as part of their PBX or ACD system, or as a standalone platform. Virtually all of those surveyed had implemented, or were considering, a form of voice processing application to improve their Call Centre's performance.

The most common form application of voice processing was the playing of messages to callers waiting in a queue. Almost half of the respondents have installed such a system, and another third were considering it.

Upfront routing of calls based on a spoken menu was also favoured by the group, with a third having such a system in place and another 40 per cent considering it. The words 'hold on the line for an agent or press 1 to order a statement, 2 for an account balance' is a single example of the customer routing themselves to device they require.

Aspen has found in the past that companies perceive interactive applications as having a higher chance of customer rejection. Proving voice processing technology with a simple application builds confidence, allowing progression to interactive systems.

Transaction processing applications were the least favoured on any application. This reflects the greater complexity of implementing such services, including integration into a company's databases.

Aspen Managing Director, Peter Massey, said: 'The survey confirms our experiences that customers and businesses alike can benefit from voice processing.

'However, an application will only succeed if you clearly understand the needs of customer segments and this is best achieved through rigorous trialling of every new service. The simpler the application, the more likely it will succeed.'

COLT Strikes New Alliances

COLT (City of London Telecommunications) has announced two important alliances; one with Energis and

one with Teleport Communications Group (TCG) in the USA.

COLT has already expanded its network westwards and is installing a second switch site.

The company says it has exceeded all expectations in the rate of growth of its network and service. Their network now covers over 70 km, with another 10 km under construction, and they provide service to over 200 buildings.

Multimedia Wars

BT's first shot in its campaign to dominate the networked multimedia market will stir up fierce competition for the telecommunications markets of the future, says leading telecoms consultants Analysys. This is because whoever wins the battle over networked multimedia services will dominate the provision of 21st century communications.

Analysys, Europe's leading telecoms consultancy, has issued a report explaining why the multimedia market will become a hard-fought battleground between telecoms operators such as BT and the cable TV companies.

According to Analysys, BT's announcement in November 1994 of market trials of video-on-demand, home shopping and banking, and multimedia educational services over the telephone network will not be unopposed.

'BT will almost certainly face a legal challenge from the cable operators in the UK,' commented Simon Norris, one of the report's authors.

'The terms of its licence currently keep it out of broadcast services by preventing it from offering entertainment services to more than one customer simultaneously. But BT will argue that on-demand services are not covered by this restriction, as customers watch when they want, when they want.'

The Analysys report gives details of the products which are available to support these services, and the plans of over 50 of the leading manufacturers and operators. It analyses the impact of networked

multimedia on the major players in the telecommunications sector, and highlights the likely emergence of a new breed of 'content provider', combining skills from the publishing, broadcast media and publishing industries.

What Users' Think

The Telecommunications Users' Association (TUA) has just published its third annual survey, which this year focuses on competition in the UK telecommunications market and, in particular, on the performance of BT and Mercury.

The survey, *Quality and Choice: Users' Perceptions of the Public Network Operators*, highlights the increasing shift away from a local service approach which the TUA believes will lead to both BT and Mercury losing revenue opportunities with small and medium sized local businesses.

The report says that local businesses are feeling increasingly distant from BT and Mercury as both companies restructure and centralise and 1995 offers a unique opportunity for new entrants operating on a local or regional level. TUA members were particularly concerned about the future of Mercury Communications itself, particularly in the light of announcements at the end of 1994.

In terms of time taken to recover after a major loss of service, both Mercury and BT performed more effectively, particularly in reducing the percentage of instances where restoration took more than two hours. According to the survey, Mercury improved the quality and performance of its network, but customers were looking for more direct access.

BT Paging received the highest rating on the performance of its service while Mercury Paging performed best on price.

Overall, customers continued to welcome the reduction in BT's prices for most services and, in particular, had welcomed the removal of peak rate, B1 routes and the reduction in directory enquiry charges.

High Speed Networks

by M. Boisseau, M. Demange, J-M. Munier

For anyone who learnt their LANs, MANs and WANs (local, metropolitan and wide area networks) in the 1980s, this is a book which will bring you up to date on what's happened since, as well as inform on a few things still in the pipeline.

After starting with a few communications basics (the open system interconnect (OSI) 7-layer model, high-level data link control (HDLC) and physical transmission media), coverage moves onto LANs. Here the workings of Ethernet and CSMA/CD (carrier sense multiple access with collision detection), token ring, token bus, FDDI (fibre distributed interface), FDDI-II, FFOL (FDDI follow-on LAN) and FCS (fibre standard channel) are described. Aspects such as access control and frame formats are well covered, as are issues such as bit-encoding and clock synchronisation. The applicability (or otherwise) of each LAN to types of application is discussed briefly.

The next topic is MANs. Although FDDI is mentioned again briefly here, the DQDB (distributed queue dual bus) protocol forms the main body of coverage in this chapter. Again frame contents and access control are well-examined.

The book now progresses onto the much larger subject of wide area networks. At the physical layer, methods of long-haul transmission (time division multiplex and frequency division multiplex) are examined and their merits compared, before closer examination of the differences between the European, American and Japanese digital transmission hierarchies. The rest of the chapter focuses on the European standards, and covers synchronous digital hierarchy (SDH) and switching. ISDN, packet switching and X.25, frame relay, cell relay (and ATM and DQDB), and broadband ISDN all receive detailed examination down to the level of frame contents and applicability. There are plenty of diagrams to show how these are interrelated, and how they relate to the OSI 7-layer model too.

Interconnection of LANs is the next topic, and sets out to examine the intricacies of bridges and routers. Transparent bridging and source routing are covered (with the spanning tree looping resolution), before moving onto routers and routing algorithms and protocols. The latter, I felt, had quite thin coverage, and could especially have done with a blow-by-blow account of how routing information actually propagates around a large internetwork.

The final chapter looks at TCP/IP (transmission control protocol/Internet protocol), examining IP addressing (although the new class D address is overlooked), the IP datagram format, ICMP (Internet control message protocol), TCP (transmission control protocol), (including connection establishment and flow control), and UDP (user datagram protocol). The OSI transport classes are covered, as is TCP's relationship with them. Finally there is a rapid look at network administration, including SNMP (simple network management protocol). This chapter covers TCP/IP only as an overview in the context of the previous chapters (volumes exist on this subject alone), so further reading would be essential to appreciate more the extent of this topic.

This book was written by three employees of IBM France, and occasionally its French origins show through the translation. This I found added an unexpected charm (rather than an irritation), which made more enjoyable the reading of a book which essentially climbs the first four layers of the OSI 7-layer model for LANs, MANs and WANs. In the light of this, I would recommend *High Speed Networks* to anyone who needs to gain a feeling of how the wide subject of networks fits together, from transmission strategies, to routing, to internetwork management.

Published by John Wiley & Sons Ltd.,
£19.95.

xii + 192pp.

ISBN 0-471-95109-9

Reviewed by Tim Hewett

book reviews

Applied Data Communications—A Business Orientated Approach

by James E. Goldman

Books on wide-ranging and rapidly evolving topics, such as data communications, tend to come in two types: either skimpy, addressing only a few technical areas; or heavyweight, resulting in an indigestible read. At over 600 pages, this book is no lightweight and gives a very comprehensive, but readable, introduction to data communications, covering almost every area of the field.

The book acknowledges its technical limitations, and tries only to increase awareness of more involved areas, without attempting a confusing explanation. References to more detailed information are extensive and, more importantly, up to date. Information is clearly presented in a down-to-earth style, which along with plenty of well-laid-out diagrams, tables and figures, make the technical

subject matter much more readable. It contains numerous questions, case studies and in-depth explanations, which allow it to be used as a tutorial as well as a reference book. Although the book is written with an American bias, there are still many up-to-date and valid references for the European reader. However, this book's unique selling point is that, unlike many similar publications, it deals with data communications from a business point of view. It positively seeks to link technical details of networks and data communications systems with customers' applications and business needs. Moreover, issues such as cost, quality, performance, upgrading legacy systems and migration strategies are introduced. Those who really do believe in putting customers first, and understanding the customers' perspective will find this approach invaluable.

The chapters on Internetworking, Enterprise Networks & Client Server Architectures, Multimedia and Voice/Data integration are particularly noteworthy. The book also attempts to

'future proof' itself with a 'Trends to Watch' chapter. In short, this is an excellent book for someone needing a very comprehensive coverage of data communications. It is equally suitable for the newcomer to the subject and those needing a useful desk reference.

Published by John Wiley & Sons Ltd.,
£51.00 (hardback), £19.95 (softback)
xxiii + 643pp.
ISBN 0-471-11084-1

Reviewed by David Tolcher

notes and comments

Erratum

In the article 'The BT Speech Applications Platform', published in the January 1995 edition of the *Journal*, the box marked 'Speech Controller' in Figure 1 on p. 305 should read 'Remote Controller'.

TELELINE ISOLATOR YOUR BEST INSURANCE



**65KV PEAK
20KV RMS**

- To protect personnel
- To protect equipment
- To maintain substation communications

DANGER

POSITRON

5101 Buchan Street
Montreal, Quebec
Canada H4P 2B9
Tel: (514) 345-2200
Fax: (514) 731-8662
1-800-361-9636

BRITISH TELECOMMUNICATIONS ENGINEERING

ISSN 0262-401X

Published in April, July, October and January by *British Telecommunications Engineering Journal*, 2-12 Gresham Street, London, EC2V 7AG. (Formerly *The Post Office Electrical Engineers' Journal* Vols. 1-74: April 1908-January 1982.)

The Board of Editors is not responsible for any statements made nor the opinions expressed in any of the articles or correspondence in this *Journal*, unless any such statement is made specifically by the Board.

© 1995: The Institution of British Telecommunications Engineers.

British Telecommunications Engineering is printed in Great Britain by Staples Printers St. Albans Ltd., Alban Park, Hatfield Road, St. Albans, Hertfordshire AL4 0JJ.

The *Structured Information Programme* is printed in Great Britain by Unwin Brothers Limited, The Gresham Press, Old Woking, Surrey, GU22 9LH.

Subscriptions and Back Numbers

Price: Single copy—companies, universities, libraries and other bodies £9.25 (£10 including postage UK, £11.25 including postage overseas); private individuals (paid by personal cheque) £5.00 (£5.75 including postage UK, £7.00 including postage overseas). Annual subscriptions (including postage)—companies etc. £40.00 UK, £45.00 overseas; private individuals (paid by personal cheque) £23.00 UK, £28.00 overseas.

Overseas customers can pay by sterling drafts drawn on London banks or by sterling Eurocheques.

Price to BT staff: £2.25 per copy.

Back numbers can be supplied if available at the single copy rates given above.

Orders, by post only, should be addressed to *British Telecommunications Engineering Journal* (Sales), PP: G012, 2-12 Gresham Street, London EC2V 7AG.

Remittances for all items (except binding) should be made payable to 'BTE Journal' and should be crossed '& Co'.

Advertisements

All enquiries relating to advertisement space reservations should be addressed to The Advertisement Manager, *British Telecommunications Engineering Journal*, PP: G012, 2-12 Gresham Street, London EC2V 7AG. (Telephone: (0171) 356 8050. Fax: (0171) 356 7942.)

Communications

All communications should be addressed to the Editorial Office, *British Telecommunications Engineering Journal*, PP: G012, 2-12 Gresham Street, London EC2V 7AG. (Telephone: (0171) 356 8050. Fax: (0171) 356 7942.)

Copyright

The entire contents of *British Telecommunications Engineering* and the *Structured Information Programme* are covered by general copyright and special permission is necessary for reprinting long extracts, but editors are welcome to use not more than one-third of any article, provided that credit is given at the beginning or end, thus: 'From British Telecommunications Engineering'.

Authorisation to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by the *British Telecommunications Engineering Journal* for users registered with the Copyright Clearance Centre's (CCC's) Transactional Reporting Service, provided that the base fee of \$5.00 per copy is paid directly to CCC, 27 Congress Street, Salem, MA 01970, USA. For those organisations that have been granted a photocopy license by CCC, a separate system of payment has been arranged. Payment is additionally required for copying of articles published prior to 1978.

0262-401X/95 \$5.00 + .00



THE INSTITUTION OF BRITISH TELECOMMUNICATIONS ENGINEERS

IBTE Local-Centre Contacts

The following is a list of Local-Centre contacts to whom enquiries about the Institution may be directed. Where no Membership Secretary details are shown, membership enquiries should be directed to the Local Secretary.

Centre	Local Secretary	Membership Secretary
East Anglia	Terry Birdseye (01702) 373330	
East Midlands	Dave Bostrom (01908) 656215	Ian Bethell (0115) 947 8587
Lancs & Cumbria	Trevor Hughes (01254) 666261	
London	Jim Skinner (0171) 728 7015	Paul Lendon 0800 306090
Manchester	David White (0161) 200 9231	Mike Ku (0161) 200 9343
Martlesham Heath	Peter Shaw (01473) 644378	Nisha Moudgill (01473) 646282
North Downs	Nick Smith (01227) 474823	
North East	Richard Jackson (0191) 261 3560	Frank Balloch (0191) 261 3226
North Wales & Marches	Peter Clay (01743) 274301	
Northern Ireland	Allan MacDougall (01232) 215301	Brenda Sharvin (01232) 320092
Scotland East	Graham Neilson (0131) 668 5878	Keith Doig (0131) 668 5315
Severnside	Mike Stenning (0117) 920 5691	Dean Shephard (0117) 920 6457
South Downs	Graham Lindsay (01273) 222633	
South Midlands	Tony Atkinson (01234) 277874	
South Wales	Peter Coleman (01222) 514524	Keith Voyce (01222) 555075
Staffordshire	Ray Piper (01785) 813483	
Thames Valley	Stephen Graham (01734) 501747	Dave Prout (01734) 501244
West Midlands	Richard Harwood (01527) 550252	Paul Jervis (0121) 230 4181
West Scotland	Graham Wellby (0141) 220 3997	Jill Maxwell (0141) 220 5837
Westward	Rob Rand (01392) 267114	Chris Gould (01392) 267113
Yorks & Lincs	Steve King (0113) 237 8124	Paul Horncastle (0113) 276 0242

Associate Section Zone Committee Contacts

The following is a list of Associate Section Zone contact points to whom enquiries about the Associate Section should be directed.

Zone Committee	Contact	Telephone No.
London	Terry McCullough	(01707) 601151
Midland	John Sansom	(0116) 253 4579
North East	Geoff Jenkinson	(0113) 237 8163
Northern Home Counties	Andy Edmonds	(01473) 645171
Northern Ireland	Brian Walker	(01232) 232413
North West	Phil Holt	(0161) 339 1971
Scotland	Raymond MacMillan	(01463) 665913
Southern Home Counties	Martyn Harvey	(01273) 568388
South West	Roger Garland	(0117) 920 5485
Wales	Ferg Brown	(01952) 249886

IBTE National HelpLine — (0171) 356 8008

Contents

Guest Editorial: Shaping BT for the Future	1
BT Networks and Systems—A Platform for Success	2
Network Measurement and Performance	5
Eliminating the Handicap of Special Needs	12
Challenges in Information Visualisation	17
When I'm 64	26
Technology and Telecommunications Beyond 2025	
Aspects of Optical Transparency	33
Global System for Mobile Communications	38
The Development of the GSM Standard	
The St. Albans Study	46
Developing Sales Engineering into Engineering Consultancy	50
Telecommunications Fraud Management	58
Reducing Costs and Winning Business	
High-Voltage Protection of Telecommunication Circuits	63
The Development of Power and Building Engineering Services in Telecommunications	71
TRIBUNE—A Broadband ISDN Test Bed	81

